Chapter 1

The First Steps

1.1. Dimension 0: Counting points using vector spaces and free abelian groups

The simplest topological spaces imaginable are just collections of finitely many points, equipped with the usual discrete topology. The only information of interest would then be the number of points in such a space S. As a foretaste of things to come, let us see how this information can be extracted in an algebraic way.

Assume the space S has n points. We now describe how to retrieve n in a way which at a first glance would appear rather *round-about*, but which illustrates what will be done in general. Let us consider a vector space over the real numbers whose basis is indexed by the elements of S. We call it $V = \mathbb{R}\langle S\rangle$, and we can formally think about points from S as the actual basis vectors in V. Now, imagine that we do not know the set S; instead we are given the vector space V without any specified basis. The question is: *what can we say about the set* S*?*

Clearly, the vector space V will have plenty of bases, and we may not find the one corresponding to S. On the other hand, a basic course in linear algebra tells us that the dimension of V is well-defined, and of course it is equal to $|S|$. So, while being unable to recover the set S itself, without further information, we can easily recover the *cardinality* of S as the dimension of the vector space that was given to us. This situation is a precursor of *homology with coefficients in the field* $\mathbb{R}$.

As the next step, note, that we could have played the same trick taking free abelian groups instead of the vector spaces. Indeed, given S, we can consider the *free abelian group* $\mathbb{Z}\langle S\rangle$ generated by the elements of S. Recall that, by definition, the elements of $\mathbb{Z}\langle S\rangle$ are all possible linear combinations $c_1\alpha_1 + \cdots + c_t\alpha_t$, where $c_1, \dots, c_t$ are integers and $\alpha_1, \dots, \alpha_t$ are elements of S, and the addition operation is the one you expect:

$$(c_1\alpha_1 + \cdots + c_t\alpha_t) + (d_1\alpha_1 + \cdots + d_t\alpha_t) = (c_1 + d_1)\alpha_1 + \cdots + (c_t + d_t)\alpha_t.$$

Again, when we are given the abelian group $\mathbb{Z}\langle S\rangle$, but not told what the set S is, we can still read off the cardinality of S from that group. This is a corollary of the classification theorem of finitely generated abelian groups: the group $\mathbb{Z}\langle S\rangle$ is free, and $|S|$ is the dimension of that free part.

Finally, we could also consider the vector space $\mathbb{Z}_2\langle S\rangle$ instead of $\mathbb{R}\langle S\rangle$. Everything becomes even easier in this case, since that vector space has finitely many points, namely $2^{|S|}$ points, so it is very easy to read off the number $|S|$ from the vector space.

Can we allow the set S to be infinite? The answer is: yes, the cardinality of S can still be determined directly from $\mathbb{R}\langle S\rangle$ or from $\mathbb{Z}_2\langle S\rangle$. However, the details are slightly more involved, and we skip them here, as we do not want to get distracted by unnecessary deviations into the realm of set theory.

It turns out that there is a wide-reaching generalization of this "counting" method, and that we have just calculated our first instance of a *homology group*! Indeed, the space $\mathbb{R}\langle S\rangle$ is called the *0th homology group of* S *with coefficients in* $\mathbb{R}$, and it is denoted $H_0(S;\mathbb{R})$. Similarly, the space $\mathbb{Z}\langle S\rangle$ is called the *0th homology group of* S *with coefficients in* $\mathbb{Z}$ or, alternatively, *with integer coefficients*, and it is denoted $H_0(S;\mathbb{Z})$. The group $H_0(S;\mathbb{Z}_2)$ is defined the same way. Let us now see what happens for slightly less trivial spaces.

1.2. Dimension 1: Graphs

Let us now move up one in dimension and consider *graphs*, which are topological spaces obtained by taking a set of vertices and then connecting them by a number of edges. For a graph G, we write $G = (V, E)$, where V denotes the set of vertices of G, and E denotes the set of edges. In order to keep things simple we assume that

- the vertices of the graph G are labeled $v_1, \dots, v_n$, with $n \geqslant 1$;
- G does not have edges glued on the same vertex with both ends (the so-called *loops*), and we also do not allow two edges to be glued on the same set of two vertices (multiple edges);
- for each edge $(v_i, v_j) \in E$ we require $i < j$.

The graphs without loops and without multiple edges are called *simple graphs*. Also our default choice is to consider graphs with only finitely many vertices and hence finitely many edges. We call these *finite graphs*.

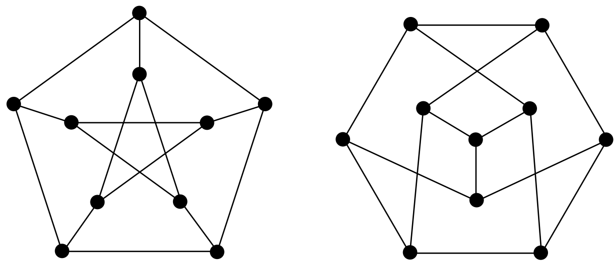

Figure 1.1. Three ways to draw the Petersen graph.

At times, graphs can confront us with surprisingly complicated problems. For example, it is hard to tell whether two graphs are isomorphic: the so-called *Graph Isomorphism Problem*; see [**KST**]. Figure 1.1 illustrates this fact by visually showing three rather different ways to draw the same graph, the so-called *Petersen graph*, whose vertices can be indexed by all 2-element subsets of a 5-element set and whose edges connect disjoint subsets.

The question of main interest in our present context is: what topological features of G can we "count", in analogy with what we did with the set S?

1.2.1. Counting connected components of a simple graph. Clearly, the first thing that we can count is the number of connected components of G. Here is how it can be done algebraically. Consider the vector space $\mathbb{R}\langle V\rangle$, with the chosen basis $v_1, \dots, v_n$. The basis vectors are in bijection with the vertices of G, and we use the same notation. For reasons which will become clear shortly, a consistent notation for that space will be $C_0(G;\mathbb{R})$, and the consistent name for this space will be the space of 0-*dimensional chains of* G *with coefficients in* $\mathbb{R}$.

Let us define an equivalence relation on the set of vertices of G by saying that *two vertices connected by an edge are equivalent.* In other words, for any two vertices $v_i, v_j \in V$, which are connected by an edge in G, that is, $(v_i, v_j) \in E$, we want to have an algebraic way of saying "set the vertex v_i equal to the vertex v_j". The standard way to do that is to consider $v_j - v_i = 0$ as a relation or, equivalently, consider a relator $v_j - v_i$. Formally, let U be the subspace of $C_0(G;\mathbb{R})$ spanned by all such relators, corresponding to edges in E, and then consider the quotient vector space $C_0(G;\mathbb{R})/U$. This quotient is again a homology group! It is denoted by $H_0(G;\mathbb{R})$, and it is called, just as above, the 0*th homology group of* G *with coefficients in* $\mathbb{R}$. We can formalize our verbal definition by writing

$$H_0(G;\mathbb{R}) := \mathbb{R}\big\langle V \,|\, v_j - v_i, \text{ for all } (v_i, v_j) \in E\big\rangle\,. \tag{1.1}$$

Let us make precise our choice of notations. Let R be some ring, let S be some set, and let T be some subset of $R\langle S\rangle$. We write $R\langle S \mid T\rangle$ to denote the quotient $R\langle S\rangle / R\langle T\rangle$. Here S is the set of generators, and T can be thought of as a set of relations. Taking the quotient is the algebraic way to say that we set certain expressions in $R\langle S\rangle$ to be equal to 0. When R is a field, the quotient $R\langle S\rangle / R\langle T\rangle$ is a vector space over R. When $R = \mathbb{Z}$, the quotient $R\langle S\rangle / R\langle T\rangle$ is an abelian group. Typically, it will not be free.

In general, the vector subspace U, defined above, is denoted by $B_0(G;\mathbb{R})$ and is called the space of *0-dimensional boundaries of G with coefficients in $\mathbb{R}$*. We leave it as an exercise, see Exercise (2)(a), to show that the dimension of the vector space $H_0(G;\mathbb{R})$ is equal to the number of connected components of G. This number is also called the *0th Betti number* of G, and it is denoted by $\beta_0(G)$.

Switching from $\mathbb{R}$ to $\mathbb{Z}_2$, the vector space $\mathbb{Z}_2\langle V\rangle$ is denoted by $C_0(G;\mathbb{Z}_2)$ and everything can be done in much the same way. We let the sums $v+w$ (which over $\mathbb{Z}_2$ is the same as $v-w$) span a vector subspace U of $C_0(G;\mathbb{Z}_2)$, and let the quotient $C_0(G;\mathbb{Z}_2)/U$ be denoted $H_0(G;\mathbb{Z}_2)$, which is then called the *0th homology group of G with coefficients in $\mathbb{Z}_2$*. In fact, taking an arbitrary field $\mathbf{k}$ we can define

$$H_0(G;\mathbf{k}) := \mathbf{k}\big\langle V \,|\, v_j - v_i, \text{ for all } (v_i, v_j) \in E\big\rangle\,.$$

Again, we leave it as an exercise, see Exercise (2b), to show that

$$\dim H_0(G;\mathbf{k}) = \beta_0(G),$$

which does not depend on the choice of the field $\mathbf{k}$.

Let us now go through the above procedure again, this time replacing the vector space $\mathbb{R}\langle V\rangle$ with the free abelian group $\mathbb{Z}\langle V\rangle$, which we denote $C_0(G;\mathbb{Z})$. We can again form a subgroup H generated by all the differences $v_j - v_i$, whenever $(v_i, v_j) \in E$. As above, consider the quotient $C_0(G;\mathbb{Z})/H$:

$$H_0(G;\mathbb{Z}) := \mathbb{Z}\big\langle V \,|\, v_j - v_i, \text{ for all } (v_i, v_j) \in E\big\rangle\,. \tag{1.2}$$

Remember that, in principle, a quotient of free abelian groups does not have to be free (though it certainly has to be abelian). Fortunately, nothing of the sort happens on the right-hand side of Equation (1.2). Exercise (2)(c) asks the reader to show that $H_0(G;\mathbb{Z})$ is free abelian, and that its dimension is equal to the Betti number $\beta_0(G)$.

Remark 1.1. When the graph G is not simple, everything can be done in almost the same way. Indeed, the space $C_0(G;\mathbb{R})$ is just the same. The loops will have no influence, since a loop based at a vertex v_i will simply give the relation $v_i - v_i = 0$ or, equivalently, will ask the vertex to be equivalent to itself. The multiple edges will also not change anything since repeating edges will just repeat the relations $v_j - v_i = 0$, which are there anyway. In other words, we can simply adopt Equation (1.1) directly.

1.2.2. Counting 1-dimensional holes in simple graphs. We would now like to understand what is going on in the context above in dimension 1. If the finite graph G is *planar*, that is, there is a way to draw it in the plane without self-intersections, then it will divide the plane into regions, with exactly one infinitely large region surrounding the graph. We could then count the number of bounded regions, and say that their number is the number of "1-dimensional holes" in the graph G. That would actually be fine, and we could let the 1*st Betti number* $\beta_1(G)$ denote that number. There is still an issue, though, with the question of whether or not this number will depend on the actual way we drew the graph, but that can be settled by a simple induction argument. In fact, it is not difficult to show, for example using induction on $|V| + |E|$, that we will have $\beta_1(G) = |E| - |V| + \beta_0(G)$, for all planar graphs G, even the non-simple ones.

A much bigger problem is that not all graphs are planar; in fact, most of them are not. If the graph is not planar, it can actually be rather confusing to talk about the number of 1-dimensional holes, or even to try to say what such a hole would actually be. In fact, it might be the sort of "intuition", which can make the understanding of what is going on harder, not easier. For example, how many holes should the complete graph on 5 vertices K_5 have? We want to say that "K_5 has 6 holes", but how can we formalize, or even visualize, that? (See Figure 1.2.)

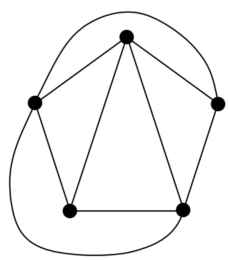

Figure 1.2. The graph K_5 without an edge has 5 holes, adding one more edge should create the 6th one.

We shall not talk about 1-dimensional holes of a graph for a simple reason: *we do not know what it means.* Instead, to draw an analogy to the previous section, we can learn how to add cycles and then measure the dimensions of the obtained groups.

Let us now describe the formal framework. Recall that the vertices of the simple graph G are labeled by $v_1, \dots, v_n$. In complete analogy to the 0-dimensional case, let $C_1(G; \mathbb{R})$ denote the vector space $\mathbb{R}\langle E \rangle$, whose basis is indexed by the edges of E. In other words, for $1 \leqslant i < j \leqslant n$, such that $(v_i, v_j) \in E$, we let e_{ij} denote the corresponding basis vector; for convenience, we shall also set $e_{ji} := -e_{ij}$, for all $1 \leqslant j < i \leqslant n$, and set $e_{ii} := 0$, for all $1 \leqslant i \leqslant n$. By construction, a vector $c \in C_1(G;\mathbb{R})$ is a linear combination $\sum_{i<j,(\nu_i,\nu_j)\in E} \alpha_{ij} e_{ij}$, where $\alpha_{ij} \in \mathbb{R}$, for all $i < j$, $(\nu_i, \nu_j) \in E$. We shall say that c is a *cycle* if for every $1 \leqslant k \leqslant n$ we have

$$\alpha_{1k} + \alpha_{2k} + \cdots + \alpha_{k-1,k} + \alpha_{k+1,k} + \cdots + \alpha_{nk} = 0, \tag{1.3}$$

where we use the handy notation $\alpha_{ji} := -\alpha_{ij}$, for $i < j$. One can see that the "old-fashioned" graph cycles are reflected in this construction as follows: a cycle $\nu_{w_1}, \nu_{w_2}, \dots, \nu_{w_t}$ corresponds to the vector

$$e_{w_1 w_2} + e_{w_2 w_3} + \cdots + e_{w_{t-1} w_t} + e_{w_t w_1}.$$

One can furthermore see that all cycles actually form a vector subspace. This can be either seen via an ad hoc computation or, more structurally, by viewing it as the kernel of a certain linear map; see the next subsection. We denote this vector space by $Z_1(G;\mathbb{R})$ and call it the *group of* 1*-cycles of* G *with real coefficients*. We shall also denote the same space $H_1(G;\mathbb{R})$ and call it the *first homology group of* G *with real coefficients*. For higher-dimensional spaces, these two groups do not have to coincide, but for graphs they do. We set the first Betti number $\beta_1(G)$ to be the dimension $\dim H_1(G;\mathbb{R})$.

As an example, let G be the complete graph on 3 vertices: $V = \{\nu_1, \nu_2, \nu_3\}$ and $E = \{(\nu_1,\nu_2),(\nu_1,\nu_3),(\nu_2,\nu_3)\}$. A cycle is a linear combination $\alpha_{12}e_{12} + \alpha_{13}e_{13} + \alpha_{23}e_{23}$ satisfying

$$\begin{cases} \alpha_{12} + \alpha_{13} = 0, \\ \alpha_{12} - \alpha_{23} = 0, \\ \alpha_{13} + \alpha_{23} = 0. \end{cases}$$

The solutions to this system form a line in $\mathbb{R}^3$ generated by the vector $(1,-1,1)$, where the ordered basis is taken to be $\{e_{12}, e_{13}, e_{23}\}$. So as the homology group we get $H_1(G;\mathbb{R}) = \mathbb{R}\langle(1,-1,1)\rangle \approx \mathbb{R}$. It makes sense to think about $e_{12} - e_{13} + e_{23}$ as a *generating cycle*.

The reader who is familiar with the fundamental group will appreciate how much easier it is to compute the first homology group of this graph as opposed to the fundamental group of the circle. On the other hand, it is not a priori clear, although it is true and easy to see in this particular case that different graph representations of the same topological space will yield isomorphic homology groups. In general, this is known as the question of *invariance under the simplicial subdivision*.

1.2.3. The boundary operator. It is useful to view what we did so far through the lens of a certain linear map

$$\partial_1 : C_1(G;\mathbb{R}) \longrightarrow C_0(G;\mathbb{R}), \tag{1.4}$$

with which we now proceed. Define the linear map ∂_1 by setting

$$\partial_1(e_{ij}) := v_j - v_i, \text{ for all } 1 \leqslant i < j \leqslant n, \tag{1.5}$$

on the basis vectors, and then extending this linearly. This map is called the *boundary operator*. We leave it as Exercise (3) to show that $B_0(G;\mathbb{R})$ is the image of ∂_1 and that $Z_1(G;\mathbb{R})$ is the kernel of this map. This yields alternative definitions for both of these groups.

The whole context can be replicated nearly verbatim, with the real coefficients replaced by integers. The kernel of the group homomorphism $\partial_1 : C_1(G;\mathbb{Z}) \longrightarrow C_0(G;\mathbb{Z})$ must itself be a free abelian group and is denoted by $H_1(G;\mathbb{Z})$.

We can also use $\mathbb{Z}_2$-coefficients instead. The kernel of the linear map $\partial_1 : C_1(G;\mathbb{Z}_2) \longrightarrow C_0(G;\mathbb{Z}_2)$ is clearly a vector space over $\mathbb{Z}_2$; it is denoted by $H_1(G;\mathbb{Z}_2)$. In fact, the case of $\mathbb{Z}_2$-coefficients is in some sense easier to deal with. In this situation, the elements of $C_1(G;\mathbb{Z}_2)$ simply correspond to graphs on n vertices. Such a graph corresponds to a cycle if and only if the valencies[1] of all vertices are even. The reader is invited to see how adding such graphs with $\mathbb{Z}_2$-arithmetic works in practice.

We remark that the identity

$$\beta_1(G) - \beta_0(G) = |E| - |V| \tag{1.6}$$

holds for all finite graphs G, not just for the planar ones. Indeed, we have

$$|E| = \dim C_1(G;\mathbb{R}) = \dim \operatorname{Ker} \partial_1 + \dim(C_1(G;\mathbb{R})/\operatorname{Ker} \partial_1), \tag{1.7}$$

$$|V| = \dim C_0(G;\mathbb{R}) = \dim \operatorname{Im} \partial_1 + \dim(C_0(G;\mathbb{R})/\operatorname{Im} \partial_1), \tag{1.8}$$

and clearly

$$\dim(C_1(G;\mathbb{R})/\operatorname{Ker} \partial_1) = \dim \operatorname{Im} \partial_1,$$

since $\partial_1 : C_1(G;\mathbb{R})/\operatorname{Ker} \partial_1 \to \operatorname{Im} \partial_1$ is an isomorphism of vector spaces. We then obtain Equation (1.6) by subtracting Equation (1.8) from Equation (1.7) and recalling that $\dim \operatorname{Ker} \partial_1 = \beta_1(G)$ and $\dim(C_0(G;\mathbb{R})/\operatorname{Im} \partial_1) = \beta_0(G)$.

1.2.4. Non-simple graphs. Again, passing on to non-simple graphs will not present much difficulty. When G is non-simple, the edges are no longer uniquely determined by their vertices. Instead, we have a set E of edges and two functions

$$\partial^\circ, \partial_\circ : E \to V,$$

where $\partial^\circ(e)$ is the initial vertex of e, and $\partial_\circ(e)$ is the terminal one. If $\partial^\circ(e) = v_i$ and $\partial_\circ(e) = v_j$, we require $i \leqslant j$. For loops we will have $\partial^\circ(e) = \partial_\circ(e)$, and for multiple edges e_1 and e_2, we will have $\partial^\circ(e_1) = \partial^\circ(e_2)$ and $\partial_\circ(e_1) = \partial_\circ(e_2)$. The vector spaces $C_1(G;\mathbb{R})$ and $C_0(G;\mathbb{R})$ are defined in the

[1] Recall that the valency of a vertex is also alternatively called the *degree* of a vertex.

same way as for the simple graphs, and we define the boundary operator $\partial_1 : C_1(G;\mathbb{R}) \to C_0(G;\mathbb{R})$ by setting

$$\partial_1(e) := \partial_\circ(e) - \partial^\circ(e), \text{ for all } e \in E.$$

Then we set $H_1(G;\mathbb{R}) := \operatorname{Ker} \partial_1$. Of course, when G is simple, this coincides with our previous definition.

1.2.5. Infinite graphs. When we have a graph $G = (V, E)$ where V or E can be infinite, everything is still the same except for one point where one has to be somewhat careful. For this, recall that elements of a vector space are *finite* linear combinations of basis vectors, even when the basis itself is allowed to be infinite. Similarly, an element of a free abelian group generated by a set X is a *finite* linear combination of the generators.

Keeping this important technical detail in mind, we can still define the boundary operator ∂_1 by taking the linear extension of Equation (1.5). Once we have the boundary operator, we simply set $H_1(G;\mathbb{R}) := \operatorname{Ker} \partial_1$ and $H_0(G;\mathbb{R}) := C_0(G;\mathbb{R})/\operatorname{Im} \partial_1$, just as before. This will extend the definition of homology to infinite graphs. In particular, the cycles will be *finite* linear combinations of edges satisfying that the *sum is* 0 *at each vertex* condition; cf. Equation (1.3).

As an example consider the infinite graph $G = (V, E)$ given by the following:

$$V := \{v_m \mid m \in \mathbb{Z}\}, \quad E := \{(v_m, v_{m+1}) \mid m \in \mathbb{Z}\}.$$

Clearly, there is no finite linear combination of edges having boundary 0, so $H_1(G;\mathbb{R}) = 0$. However, there exists an infinite linear combination of edges having boundary 0, namely take the sum of *all* edges $\sum_m e_{m,m+1}$. Following up on this innocent looking example will eventually lead to *Borel-Moore homology*; see [**BM60**, **Br97**].

1.3. Dimension 2: Losing the freedom

1.3.1. Simplicial complexes of dimension 2. To proceed with dimension 2 we need a slightly more formal definition than the one we have had for graphs. The 2-dimensional simplicial complexes can be obtained from simple graphs by filling out some of the triangles. In order to keep technicalities as simple as possible, we restrict our framework a little bit further by considering the so-called plain complexes.

Definition 1.2. A *plain 2-dimensional simplicial complex* K consists of a finite simple graph $G = (V, E)$, together with a set of triangles T, where

(1) each triangle is a triple (v_i, v_j, v_k), for some $v_i, v_j, v_k \in V$, such that $i < j < k$;

(2) if $(v_i, v_j, v_k) \in T$, then all the 3 tuples (one should think of them as edges) (v_i, v_j), (v_i, v_k), and (v_j, v_k) must also lie in E;

(3) all triples in T are distinct.

We shall set $K(0) := V$, $K(1) := E$, and $K(2) := T$.

1.3.2. Homology groups in dimensions 0 and 2. Let us set

$$C_0(K;\mathbb{R}) := \mathbb{R}\langle V\rangle\,, \quad C_1(K;\mathbb{R}) := \mathbb{R}\langle E\rangle\,, \text{ and } C_2(K;\mathbb{R}) := \mathbb{R}\langle T\rangle\,.$$

The 0th homology group $H_0(K;\mathbb{R})$ is the same as $H_0(G;\mathbb{R})$; to define it, we simply ignore that we have added triangles.

Let us describe the second homology group $H_2(K;\mathbb{R})$. First, for all $1 \leqslant i < j < k \leqslant n$, such that $(v_i, v_j, v_k) \in T$, we let t_{ijk} denote the corresponding generator of $C_2(K;\mathbb{R})$. For convenience, we set

$$\begin{aligned} t_{jki} &:= t_{kij} := t_{ijk} \text{ and} \\ t_{jik} &:= t_{ikj} := t_{kji} := -t_{ijk}. \end{aligned} \tag{1.9}$$

An easy way to remember this sign rule is to say that each time some two neighboring indices are swapped, the generator changes its sign.

A 2-*chain* is an expression of the type $\sum_{i,j,k} \alpha_{ijk} t_{ijk}$, where $\alpha_{ijk} \in \mathbb{R}$, and the summation is taken over all triples (i,j,k), such that $i < j < k$, and $(v_i, v_j, v_k) \in T$. We say that a 2-chain is a 2-*cycle* if for any $(v_i, v_j) \in E$ we have

$$\sum_{k \neq i,j} \alpha_{ijk} = 0, \tag{1.10}$$

where, in analogy to Equations (1.3) and (1.9) we use the notation $\alpha_{jki} := \alpha_{kij} := \alpha_{ijk}$ and $\alpha_{jik} := \alpha_{ikj} := \alpha_{kji} := -\alpha_{ijk}$. Note that over $\mathbb{Z}_2$, Equation (1.10) just says that we have selected a set of triangles such that each edge belongs to an even number of these triangles.

Now the groups $Z_2(K;\mathbb{R})$ and $H_2(K;\mathbb{R})$ are both set to be the group of all 2-cycles. Again, one can show that the latter is actually a vector subspace. This can be done either by an ad hoc computation or by viewing it as the kernel of the linear map defined in the next subsection.

Completely analogously, one can define $H_2(K;\mathbb{Z})$ and $H_2(K;\mathbb{Z}_2)$. In Exercise (4) the reader is asked to show that $H_2(K;\mathbb{Z})$ is a free abelian group.

1.3.3. Homology group in dimension 1. To start with, the group of 1-cycles $Z_1(K;\mathbb{R})$ is the same as $H_1(G;\mathbb{R})$. As the second ingredient for defining the first homology group, for each $i < j < k$ such that $(v_i, v_j, v_k) \in T$ we want to introduce a relation

$$e_{ij} + e_{jk} = e_{ik}.$$

The way to do it is to span the subspace, which we call $B_1(K;\mathbb{R})$, by all the elements $e_{ij} - e_{ik} + e_{jk}$, for $i < j < k$, $(v_i, v_j, v_k) \in T$; note that each such element can alternatively be written in the cyclic fashion as $e_{ij} + e_{jk} + e_{ki}$. Then take the quotient $Z_1(K;\mathbb{R})/B_1(K;\mathbb{R})$. The quotient is well-defined, since any element $e_{ij} - e_{ik} + e_{jk}$ is obviously a cycle. Taking this quotient is an algebraic way of saying that we allow the cycles to be deformed across triangles, thus modeling the continuous notion of homotopy. Per definition, we now set $H_1(K;\mathbb{R}) := Z_1(K;\mathbb{R})/B_1(K;\mathbb{R})$.

Again, it is rather useful to phrase everything we are doing using linear maps as the so-called *boundary operators*. In fact, here we have two such maps: one map in dimension 1,

$$\partial_1 : C_1(K;\mathbb{R}) \longrightarrow C_0(K;\mathbb{R}),$$

defined by $\partial_1(e_{ij}) = v_j - v_i$, for all $i < j$; and one map in dimension 2,

$$\partial_2 : C_2(K;\mathbb{R}) \longrightarrow C_1(K;\mathbb{R}),$$

defined by $\partial_1(t_{ijk}) = e_{ij} - e_{ik} + e_{jk}$, for all $i < j < k$. We can then set

$$Z_1(K;\mathbb{R}) := \operatorname{Ker} \partial_1, \quad Z_2(K;\mathbb{R}) := \operatorname{Ker} \partial_2,$$

$$B_0(K;\mathbb{R}) := \operatorname{Im} \partial_1, \text{ and } B_1(K;\mathbb{R}) := \operatorname{Im} \partial_2.$$

To make notations uniform, we furthermore set

$$Z_0(K;\mathbb{R}) := C_0(K;\mathbb{R}) \text{ and } B_2(K;\mathbb{R}) := 0.$$

Then we can summarize what we have defined so far as

$$H_0(K;\mathbb{R}) = Z_0(K;\mathbb{R})/B_0(K;\mathbb{R}),$$
$$H_1(K;\mathbb{R}) = Z_1(K;\mathbb{R})/B_1(K;\mathbb{R}),$$
$$H_2(K;\mathbb{R}) = Z_2(K;\mathbb{R})/B_2(K;\mathbb{R}).$$

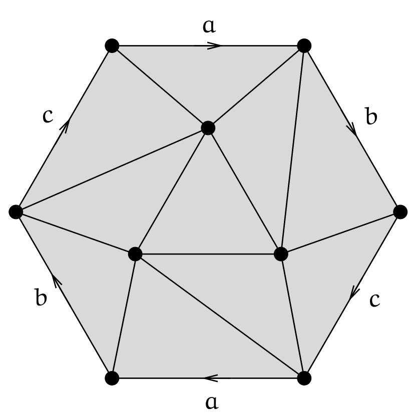

Figure 1.3. Simplicial complex X with torsion in $H_1(X;\mathbb{Z})$. This is a triangulation of the real projective plane.

Let us now see what happens here when we use integer coefficients instead. Clearly, $Z_1(K;\mathbb{Z})$ is a free abelian group, since it is the kernel of a group homomorphism from a free abelian group, and $B_1(K,\mathbb{Z})$ is a free abelian group, since it is the image of a group homomorphism into a free abelian group. However, as we mentioned before, when one takes a quotient of two free abelian groups, it can happen that the quotient is not free. We have said it already, but here it is the first time that we have the situation where this actually might happen. In this case, one speaks of the existence of *torsion* in the corresponding homology group. For torsion to occur, we need to have a 1-cycle c which is not a 1-boundary, but for which there exists a positive integer m such that $m \cdot c$ is a 1-boundary. Figure 1.3 shows an example of such a situation.

Exercises

(1) The maximal connected subgraphs of G are called *the connected components*. Show by elementary methods that for any graph G without self-crossings on the 2-dimensional sphere we have the formula

$$v + r = e + c + 1,$$

where v is the number of vertices, e is the number of edges, c is the number of connected components of G, and r is the number of regions into which the sphere is divided by our graph.

(2) Let G be a finite simple graph. Let c denote the number of connected components of G.
   (a) Show that $H_0(G;\mathbb{R}) \approx \mathbb{R}^c$.
   (b) Show that $H_0(G;\mathbf{k}) \approx \mathbf{k}^c$, where $\mathbf{k}$ is an arbitrary field.
   (c) Show that $H_0(G;\mathbb{Z})$ is a free abelian group and that furthermore we have an isomorphism $H_0(G;\mathbb{Z}) \approx \mathbb{Z}^c$.

(3) Let G be a finite simple graph. Show that $B_0(G;\mathbb{R})$ is the image of ∂_1 and that $Z_1(G;\mathbb{R})$ is the kernel of the boundary map defined by Equation (1.5).

(4) Assume K is a plain 2-dimensional simplicial complex. Show that $H_2(K;\mathbb{Z})$ is a free abelian group.

(5) Construct a plain connected[2] 2-dimensional simplicial complex K such that
   (a) $H_1(K;\mathbb{Z}) \approx \mathbb{Z} \oplus \mathbb{Z}_2$;
   (b) $H_1(K;\mathbb{Z}) \approx \mathbb{Z}_3$.

[2]Such a complex is called connected if the underlying edge graph is connected.

(6) If you are familiar with the concept of simplicial subdivision, construct a plain 2-dimensional simplicial complex K such that its second homology group is non-trivial, i.e., $\mathsf{H}_2(\mathsf{K};\mathbb{R}) \neq 0$, yet it does not contain a simplicial subdivision of a 2-dimensional sphere. What is the minimal number of triangles that one needs?

Organizing Collapsing Sequences

10.1. Face poset of an abstract simplicial complex

In order to learn how to keep track of various collapsing sequences, it is useful to introduce some combinatorial notions which encode the simplicial structure.

Definition 10.1. Let $\mathcal{K}$ be an arbitrary abstract simplicial complex. The *face poset* of $\mathcal{K}$ is a partially ordered set, which we denote by $\mathsf{F}(\mathcal{K})$, defined by the following:

- the elements of $\mathsf{F}(\mathcal{K})$ are all the simplices of $\mathcal{K}$, including the empty one;
- the partial order is given by the inclusion relation on the simplices, in other words, we set $\sigma \geqslant \tau$ as elements of $\mathsf{F}(\mathcal{K})$ if and only if $\sigma \supseteq \tau$ as simplices.

Note that sometimes one defines the face poset as the partially ordered set of all *non-empty* simplices. There are advantages to both conventions. When used to record collapses, it is handy to have the empty simplex included.

As an example, when $\mathcal{K}$ is the n-simplex, whose set of vertices is $[n]$, we obtain the partially ordered set consisting of all the subsets of $[n]$. This partially ordered set is called the *Boolean algebra*.

The next definition describes a procedure which in some sense is the reverse of taking the face poset. For this, recall that in a partially ordered

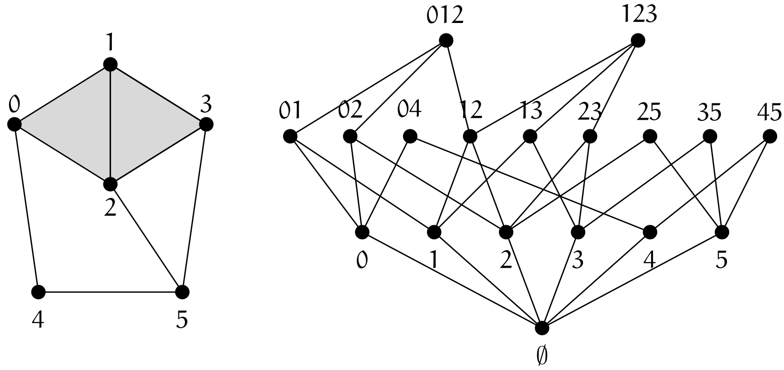

Figure 10.1. An abstract simplicial complex and its face poset.

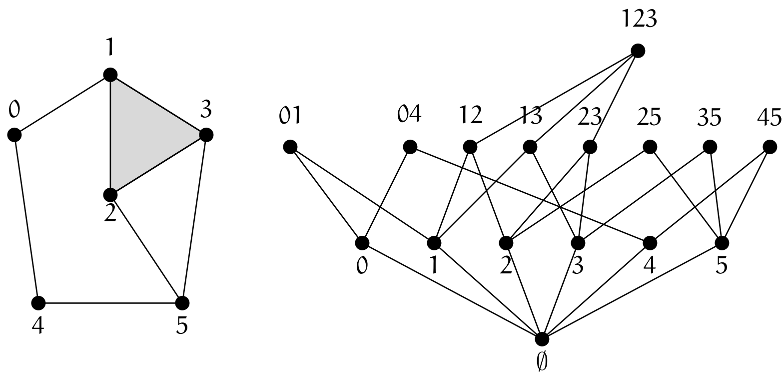

Figure 10.2. The effect of an elementary simplicial collapse $(\{0,1,2\},\{0,2\})$ on the face poset of the corresponding abstract simplicial complex.

set, a *chain* is a subset of elements which is totally ordered. In particular, a set consisting of a single element of P is a chain, and so is the empty set.

Definition 10.2. Let P be a partially ordered set. Its *order complex* is an abstract simplicial complex, which we denote by $\Delta(\mathsf{P})$, defined as follows:

- we take the elements of P as the vertices of $\Delta(\mathsf{P})$;
- the simplices of $\Delta(\mathsf{P})$ are precisely all the finite chains of P.

The next proposition describes the precise manner in which the constructions described in Definitions 10.1 and 10.2 interact. Before that, just a piece of notation: for any abstract simplicial complex $\mathcal{K}$, we denote the minimal element of $\mathsf{F}(\mathcal{K})$ by $\hat{0}$; it corresponds to the empty simplex.

Proposition 10.3.

(1) *Assume* $\mathcal{K}$ *is an abstract simplicial complex; then* $\Delta(\mathsf{F}(\mathcal{K}) \setminus \{\hat{0}\})$ *is the barycentric subdivision of* $\mathcal{K}$.

(2) *Assume* P *is a partially ordered set; then the poset* $\mathsf{F}((\Delta(\mathsf{P}))$ *is the poset of chains of* P, *including the empty one.*

Proof. Assume first that we have an abstract simplicial complex $\mathcal{K}$. Unwinding the definitions we see that the non-empty simplices of $\Delta(\mathsf{F}(\mathcal{K}) \setminus \{\hat{0}\})$ are all the totally ordered sets of non-empty simplices of $\mathcal{K}$. These, of course, are exactly the chains of $\mathsf{F}(\mathcal{K}) \setminus \{\hat{0}\}$, so comparing this with the description of the simplicial structure given in Definition 2.41 we arrive at the desired conclusion.

The second part involving the partially ordered set P is immediate as well, once the definitions of F and Δ have been unwinded. □

Most of the combinatorial constructions involving abstract simplicial complexes have their poset interpretation. Figure 10.2 shows the effect of an elementary simplicial collapse. Figures 10.3 and 10.4 show the meaning of the deletion and the link in the face poset.

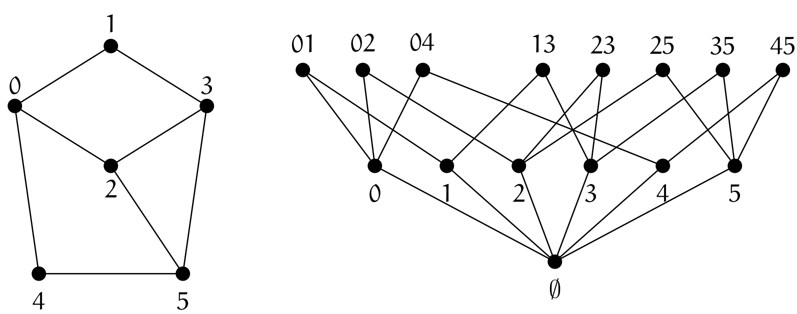

Figure 10.3. Deletion of the simplex $\{1, 2\}$ from the simplicial complex in Figure 10.1 and the corresponding face poset.

Figure 10.5 illustrates the case of the stellar subdivision. When passing from the face poset of an abstract simplicial complex to the face poset of its stellar subdivision one needs to perform the so-called *combinatorial blowup*. Before we proceed with the definition, recall the following terminology.

Let P be a poset, and choose a subset $A \subseteq \mathsf{P}$. Consider the set S of all *lower bounds* for A, i.e., $S := \{z \mid z \leqslant x,\ \forall x \in A\}$. If the set S has a unique maximal element, then it is called the *meet* of A, and is denoted by $\bigwedge_{\mathsf{P}} A$, or simply $\bigwedge A$. The poset P is called a *meet-semilattice*, or simply a *semilattice*, if it has a meet for any non-empty finite subset.

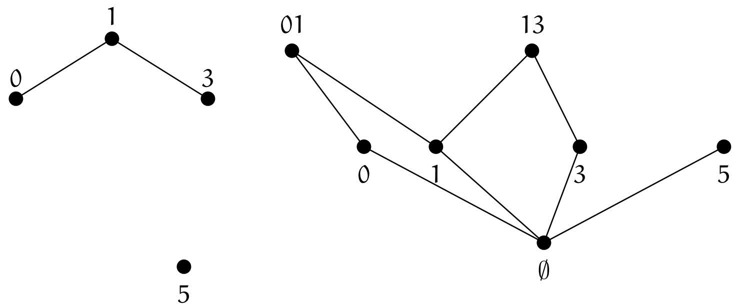

Figure 10.4. Link of the simplex {2} in the simplicial complex in Figure 10.1 and the corresponding face poset.

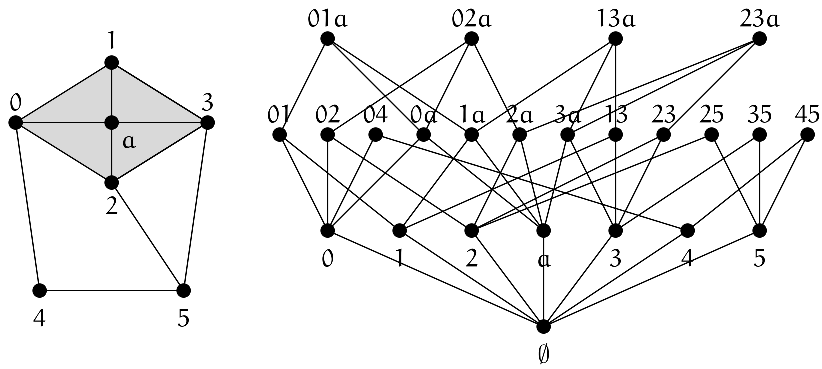

Figure 10.5. Stellar subdivision of the simplex $\{1, 2\}$ in the simplicial complex in Figure 10.1 and the corresponding face poset.

When A consists of two elements, $A = \{x, y\}$, the meet of A is called the meet of x and y, and is denoted by $x \wedge_P y$, or simply by $x \wedge y$.

Dually, for a subset $A \subseteq P$. Consider the set S of all *upper bounds* for A, i.e., $S := \{z \,|\, z \geqslant x,\ \forall x \in A\}$. If the set S has a unique minimal element, then it is called the *join* of A, and is denoted by $\vee_P A$, or simply $\vee A$. For $A = \{x, y\}$ we write $x \vee_P y$, or simply $x \vee y$.

Definition 10.4. For a semilattice L and an element $a \in L$, $a \neq \hat{0}$,[1] we define a poset $\mathrm{Bl}_a L$, the *combinatorial blowup* of L at a, as follows. The elements of $\mathrm{Bl}_a L$ are given by

(1) $y \in L$, such that $y \not\geqslant a$;
(2) $\langle a, y\rangle$, for $y \in L$, such that $y \not\geqslant a$ and $y \vee_L a$ exists (in particular, $\left\langle a, \hat{0}\right\rangle$ can be thought of as the result of blowing up a).

The order relations in $\mathrm{Bl}_a L$ are given by

[1] Here, $\hat{0}$ denotes the unique minimal element of L.

(1) $y > z$ in $\mathrm{Bl}_a L$ if $y > z$ in L;
(2) $\langle a, y\rangle > \langle a, z\rangle$ in $\mathrm{Bl}_a L$ if $y > z$ in L;
(3) $\langle a, y\rangle > z$ in $\mathrm{Bl}_a L$ if $y \geqslant z$ in L,

where in all three cases $y, z \not\geqslant a$.

As a special case, we can easily describe the combinatorics of the simplicial complex obtained by taking an n-simplex, and taking the stellar subdivision of one of its simplices. Let $\mathcal{K}$ denote the n-simplex whose set of vertices is $[n]$. Pick $S \subseteq [n]$, and let σ be the corresponding simplex of $\mathcal{K}$. Then, the set of all the simplices of $\mathrm{Sd}_{\mathcal{K}}(\sigma)$ is given by

$$\{T \mid T \not\supseteq S\} \cup \{\langle a, T\rangle \mid T \not\supseteq S\},$$

where a is the symbol corresponding to the barycenter of σ. Note that the actual new vertex at the barycenter of σ is denoted by $\langle a, \emptyset\rangle$. The simplex inclusion rules are then simply

(1) the simplex indexed by T_1 contains the simplex indexed by T_2 if and only if $T_1 \supseteq T_2$;
(2) $\langle a, T_1\rangle$ contains $\langle a, T_2\rangle$ if and only if $T_1 \supseteq T_2$;
(3) $\langle a, T_1\rangle$ strictly contains T_2 if and only if $T_1 \supseteq T_2$,

where in all three cases $T_1, T_2 \not\supseteq S$.

10.2. Acyclic matchings

Matching is the combinatorial notion for the face posets corresponding to elementary simplicial collapses. Let us recall this concept from graph theory.

Definition 10.5. Let G be a graph whose set of vertices is V and whose set of edges is E. A *partial matching* in G is a set of edges $\{\{a_1, b_1\}, \dots, \{a_t, b_t\}\}$ such that all the vertices $\{a_1, \dots, a_t, b_1, \dots, b_t\}$ are distinct.

Our notion of the partial matching is flexible, in the sense that it includes the case where *all* the vertices are matched. This makes writing arguments easier. However, often it is useful to specifically point out that all the vertices have really been matched, in which case we may also call such a matching *a complete matching.* In this text, we shall never drop the adjective "complete", so when we simply say "matching", we shall always mean the partial matching.

It is often convenient to think of a matching in a formal way: namely, as a function $\mu : M \to M$, where M is a subset of the set of vertices of G. This function must satisfy the following two conditions:

- for all $v \in M$, the vertices v and $\mu(v)$ are connected by an edge, called the *matching edge*;

- for all $v \in M$, we have $\mu(\mu(v)) = v$.

The correspondence with the matchings is easy: M is the set of matched vertices, and each vertex $v \in M$ is matched to the vertex $\mu(v)$.

Matching theory is an extensive branch of graph theory, with many methods developed to find new matchings and to improve existing ones. We refer the reader to [**Lo86**] as a possible point of entry.

In general, there are many constructions which associate a graph to a poset P. The one we need here takes the set of vertices of P as the set of vertices of that graph, and then connects two elements by an edge if and only if one of these elements covers the other one. To this end, recall that for $x, y \in P$ we say that x *covers* y, and write $x \succ y$, if $x > y$ and there exists no $z \in P$, such that $x > z > y$. The obtained graph is called the *underlying graph of the Hasse diagram of* P.

Definition 10.6. A *partial matching* in a poset P is a partial matching in the underlying graph of the Hasse diagram of P. In other words, it is a subset $M \subseteq P$, together with a bijection $\mu : M \to M$, such that for all $v \in M$, the following two conditions are satisfied:

- either v covers $\mu(v)$, or vice versa;
- $\mu(\mu(v)) = v$.

We shall think about the set M as being a part of the information provided by the the function μ, so we shall simply say things like "assume we have a matching μ."

Given a bijection $\mu : M \to M$ as in Definition 10.6, for future reference, we introduce the following notation:

- $M^{\uparrow}$ is the subset of all $v \in M$, such that $\mu(v)$ is covered by v;
- $M^{\downarrow}$ is the subset of all $v \in M$, such that $\mu(v)$ covers v;
- $R(\mu)$ is the complement of M, i.e., $R(\mu) := P \setminus M$.

Note that $M^{\uparrow} \cup M^{\downarrow} = M$, the union is disjoint, and the maps $\mu : M^{\uparrow} \to M^{\downarrow}$ and $\mu : M^{\downarrow} \to M^{\uparrow}$ are well-defined bijections, which are inverses of each other. When μ is clear, we shall simply write R instead of $R(\mu)$.

As a more precise piece of notation, we write $\mu_-(v)$ instead of just $\mu(v)$ if v covers $\mu(v)$. Symmetrically, we write $\mu_+(v)$ instead of $\mu(v)$ if $\mu(v)$ covers v.

An elementary simplicial collapse in an abstract simplicial complex $\mathcal{K}$ is now encoded as the matching of two vertices in the face poset $F(\mathcal{K})$, subject to further conditions. These vertices correspond to the two simplices σ and τ which are removed during the collapse, and the fact that they are connected by an edge is ensured by the conditions $\sigma \supset \tau$ and $\dim \tau + 1 = \dim \sigma$.

Accordingly, a set of elementary simplicial collapses is described by a matching consisting of a collection of pairs of simplices (σ, τ), such that σ contains τ, and $\dim \sigma = \dim \tau + 1$. It is a simple, but crucial observation, that not every matching of this type can be turned into a collapsing sequence. For example, no order can be chosen in the matching on the right of Figure 10.6, which would correspond to an allowed collapsing sequence.

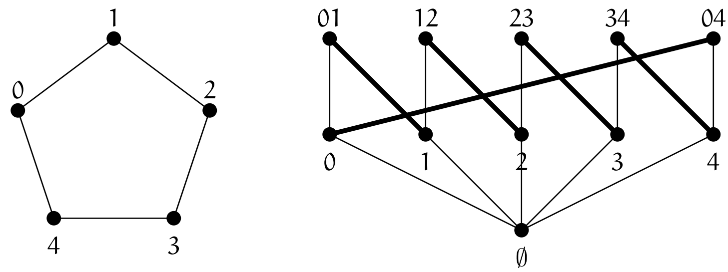

Figure 10.6. A cycle and a matching in its face poset.

Here is what goes wrong in this example: the prospective collapses are all "hooked up" with each other in a cyclic pattern, so that any suggested collapsing sequence would have one of these edges from the cycle occurring before the other ones, clearly contradicting the conditions for being an elementary collapse. This simple observation leads to the following formalization.

Definition 10.7. Assume we are given a partially ordered set P, and a partial matching $\mu : M \to M$ on P. This matching is called *acyclic* if there does not exist a cycle of the following form:

$$b_1 \succ \mu(b_1) \prec b_2 \succ \mu(b_2) \prec \cdots \prec b_n \succ \mu(b_n) \prec b_1, \tag{10.1}$$

with $n \geqslant 2$ and all $b_i \in M$ being distinct.

A graphic way to reformulate condition (10.1) of Definition 10.7 is as follows. Given a poset P, we start by orienting all edges in the underlying graph of the Hasse diagram of P, so that each one points from the larger element to the smaller one. This graph is obviously acyclic.[2] Now, assume we are given a partial matching $\mu : M \to M$. For each $v \in M$, such that v is covered by $\mu(v)$, change the orientation of the edge $(\mu(v), v)$ to the opposite one. The condition in question now says that the directed graph, obtained in this fashion, has no cycles.

Let us next formulate a proposition which provides yet another alternative reformulation of Definition 10.7. Consider first the following general construction. Assume P is a poset, and the element set of P is partitioned

[2]A graph is called *acyclic* if it does not have any cycles.

into non-empty disjoint sets $\{A_i\}_{i\in Q}$. We define a partial order on the index set Q using the following two rules:

(1) for $i, j \in Q$, we write $i \succ j$ if there exist $x \in A_i$ and $y \in A_j$, such that $x > y$;

(2) for $i, j \in Q$, we write $i >_Q j$ if there exists a finite sequence $i_1, \dots, i_t \in Q$, $t \geqslant 2$, such that $i_1 = i$, $i_t = j$, and $i_k \succ i_{k+1}$, for all $k = 1, \dots, t-1$.

The resulting partial order on Q is said to be *induced by* P. It may or may not be well-defined.

Proposition 10.8. *Assume* P *is a poset equipped with a partial matching* μ. *Let* $\{A_i\}_{i\in Q}$ *be the partition of* P *induced by that matching, where matched pairs of elements form 2-sets and non-matched elements give singletons.*

The matching μ *is acyclic if and only if the partial order on* Q *induced by* P *is well-defined.*

The proof of Proposition 10.8 is straightforward and is best left to the interested reader. From the point of view of universal constructions the thinking along the lines of Proposition 10.8 is rather fruitful. We will return to this topic in the context of the universality of the colimit of a matching in Theorem 16.5, where more details will be provided.

10.3. Collapsing sequences vs acyclic matchings: Theorem A

The next theorem is the first, and the simplest, of the central results of discrete Morse theory. In essence, it states that acyclic matchings provide a perfect language for saying that one abstract simplicial complex collapses to another one.

Theorem 10.9 (Theorem A). *Assume* $\mathcal{K}$ *is an abstract simplicial complex, and assume* $\mathcal{K}'$ *is a simplicial subcomplex of* $\mathcal{K}$, *such that* $\mathcal{K} \setminus \mathcal{K}'$ *is finite. The following statements are equivalent:*

(1) *there exists a sequence of elementary collapses leading from* $\mathcal{K}$ *to* $\mathcal{K}'$;

(2) *there exists a complete acyclic matching*[3] *on the set of all simplices of* $\mathcal{K}$ *which are not contained in* $\mathcal{K}'$.

Proof. Let us first show that (1) implies (2). Fix some sequence of elementary simplicial collapses leading from $\mathcal{K}$ to $\mathcal{K}'$, and take the matching μ on the set of simplices of $\mathcal{K} \setminus \mathcal{K}'$, which corresponds to this sequence.

[3] By this we mean *matching on the underlying graph of the Hasse diagram of the face poset of* $\mathcal{K}$.

Assume that this matching is not acyclic. Then, by definition, there must exist a cycle of the form

$$b_1 \succ \mu(b_1) \prec b_2 \succ \mu(b_2) \prec \cdots \prec b_n \succ \mu(b_n) \prec b_1,$$

for some distinct elements $b_1, \dots, b_n \in \mathcal{K} \setminus \mathcal{K}'$.

Consider the sequence $(b_1, \mu(b_1)), \dots, (b_n, \mu(b_n))$ of elementary collapses. Without loss of generality, we can assume that in the sequence of elementary collapses leading from $\mathcal{K}$ to $\mathcal{K}'$, the elementary collapse $(b_1, \mu(b_1))$ occurs before all the other elementary collapses from this set. Clearly, this contradicts the fact that $\mu(b_1) \prec b_2$, since then $\mu(b_1)$ is properly contained in at least two simplices, namely in b_1 and in b_2, both of which are present in the complex when we attempt to perform the elementary collapse $(b_1, \mu(b_1))$.

Let us now show the reverse direction, that is, that (2) implies (1). Consider a complete acyclic matching on the set of simplices $\mathcal{K}\setminus\mathcal{K}'$. We shall show that there exists a collapsing sequence from $\mathcal{K}$ to $\mathcal{K}'$ using induction on $|\mathcal{K} \setminus \mathcal{K}'|$. When $|\mathcal{K} \setminus \mathcal{K}'| = 1$, the statement is trivial, and we take it as the basis for induction.

For the induction step, assume $|\mathcal{K} \setminus \mathcal{K}'| \geqslant 2$. We construct a directed graph G as follows. As vertices of G we take the elements of $\mathcal{K} \setminus \mathcal{K}'$. We let edges be defined by the containment relations, that is, (σ, τ) is an edge if $\sigma \supset \tau$ as simplices. Finally, whenever $\sigma \supset \tau$, we orient the edge (σ, τ) *from* σ *to* τ, unless $\tau = \mu_-(\sigma)$, in which case we orient the edge *from* τ *to* σ instead. As already mentioned above, the fact that the matching is acyclic is equivalent to saying that the obtained directed graph is acyclic, i.e., does not have any *oriented* cycles.

It is a standard fact of graph theory that an acyclic directed graph contains at least one *source*, which is a vertex v such all the edges adjacent to v are oriented away from v. Let $\tau \in \mathcal{K}\setminus\mathcal{K}'$ denote such a source. There cannot exist $\sigma \in \mathcal{K} \setminus \mathcal{K}'$ such that $\tau = \mu_+(\sigma)$, because then the edge (σ, τ) would be oriented from σ to τ, contradicting the assumption that τ is a source. Since the matching is complete, we must instead have some $\sigma \in \mathcal{K} \setminus \mathcal{K}'$, such that $\tau = \mu_-(\sigma)$.

We now claim that (σ, τ) corresponds to an elementary simplicial collapse in $\mathcal{K}$. Indeed, since τ is a source, there exists no vertex γ with an edge towards τ, which means that there exists no simplex $\gamma \neq \sigma$ which properly contains τ. On the other hand, the simplex τ is of course a boundary simplex of σ of codimension 1, since these two simplices are matched, thus yielding a contradiction.

Set $\widetilde{\mathcal{K}} := \mathcal{K} \setminus \{\sigma, \tau\}$. Clearly, the simplicial complex $\mathcal{K}'$ is a subcomplex of $\widetilde{\mathcal{K}}$, since we have removed two simplices outside of $\mathcal{K}'$. The acyclic matching above gives a complete acyclic matching on $\widetilde{\mathcal{K}} \setminus \mathcal{K}'$. This matching has one edge less, so by induction hypothesis, there exists a collapsing sequence from $\widetilde{\mathcal{K}}$ to $\mathcal{K}'$. Concatenating this sequence with the elementary collapse (σ, τ) will give the desired collapsing sequence from $\mathcal{K}$ to $\mathcal{K}'$. □

10.4. Collapsing sequences and cones

10.4.1. Canonical way to collapse cones. Let $\mathcal{K}$ be a finite abstract simplicial complex, and let $\mathcal{L}$ be a cone over $\mathcal{K}$ with an apex a. There is a canonical way to collapse this cone. Before getting into the description of the collapsing procedure, we introduce a notation which will come in handy in many situations.

Definition 10.10. Assume we have a set A, a subset $B \subseteq A$, and an element $v \in A$; then we set

$$B \,\mathrm{XOR}\, v := \begin{cases} B \cup v, & \text{if } v \notin B; \\ B \setminus v, & \text{otherwise.} \end{cases}$$

The operation XOR is called *exclusive OR*.

We note the following useful identity:

$$(B \,\mathrm{XOR}\, v) \,\mathrm{XOR}\, v = B, \tag{10.2}$$

which is valid for all v and B. We also have $|B \,\mathrm{XOR}\, v| = |B| \pm 1$.

Any cone can now be collapsed by taking the exclusive OR with its apex.

Proposition 10.11. *The abstract simplicial complex $\mathcal{L}$ described above is collapsible.*

Proof. Consider the matching $\mu : \mathcal{L} \to \mathcal{L}$ on the set of simplices of $\mathcal{L}$, given by the rule

$$\mu(\sigma) := \sigma \,\mathrm{XOR}\, a.$$

Since the vertex a can be added to any simplex $\sigma \in \mathcal{K}$, the function μ is well-defined, and Equation (10.2) tells us that $\mu(\mu(\sigma)) = \sigma$. Furthermore, for all σ, $\dim \mu(\sigma) = \dim \sigma \pm 1$, and either $\mu(\sigma)$ is contained in σ or vice versa. This means that μ is actually a matching. It is also clearly complete.

Assume now μ is not acyclic, and take a cycle

$$b_1 \succ \mu(b_1) \prec b_2 \succ \mu(b_2) \prec \cdots \prec b_n \succ \mu(b_n) \prec b_1,$$

where $n \geqslant 2$, and all the simplices are distinct. By definition of μ we have $a \in b_1$, $a \in b_2$, and $a \notin \mu(b_1)$. This means $b_1 = \mu(b_1) \cup a = b_2$, yielding a contradiction. □

10.4.2. Stellar and barycentric subdivisions of a simplex. Let n be a positive integer, and let $\mathcal{K}$ be an n-simplex. Let $\mathcal{L}$ be the stellar subdivision of $\mathcal{K}$ with respect to the top-dimensional simplex. By definition, the set of vertices of $\mathcal{L}$ is $\{v_1, \dots, v_{n+1}, a\}$, where $\{v_1, \dots, v_{n+1}\}$ are the vertices of $\mathcal{K}$, and a is the new vertex added at the barycenter of the top-dimensional simplex. The simplices of $\mathcal{L}$ are given by the set

$$\{\sigma \,|\, \sigma \in \mathcal{K},\ \sigma \neq (v_1, \dots, v_{n+1})\} \cup \{a \cup \sigma \,|\, \sigma \in \mathcal{K},\ \sigma \neq (v_1, \dots, v_{n+1})\},$$

where in both sets σ is allowed to be empty.

It is obvious from this description that $\mathcal{L}$ can be obtained by starting from the $(n+1)$-simplex, and then performing an elementary collapse.

The complex $\mathcal{L}$ is a cone over $\mathcal{K} \setminus (v_1, \dots, v_{n+1})$ with apex at a, so it follows from Proposition 10.11 that $\mathcal{L}$ is collapsible.

As another example, we can take the barycentric subdivision of the simplex $\mathcal{K}$. We have seen that this subdivision is a cone with apex at the vertex indexed with the maximal simplex of $\mathcal{K}$, hence, again by Proposition 10.11 we know that it is collapsible.

10.4.3. Removing a simplex whose link is collapsible. One of the first things one can do to simplify the abstract simplicial complex at hand is to remove all the simplices whose links are collapsible themselves. As the first simple, but instructive, case, let us see how to remove a simplex whose link is actually a cone.

Proposition 10.12. *Assume $\mathcal{K}$ is a finite abstract simplicial complex, and assume σ is a simplex of $\mathcal{K}$, such that* $\mathrm{lk}_{\mathcal{K}}(\sigma)$ *is a cone. Then, we have* $\mathcal{K} \searrow \mathrm{dl}_{\mathcal{K}}(\sigma)$.

Proof. The argument is a modification of the proof of Proposition 10.11. Let b denote the apex of the cone $\mathrm{lk}_{\mathcal{K}}(\sigma)$. For an arbitrary simplex τ, which lies in the open star of σ, we set

$$\mu(\tau) := \tau \,\mathrm{XOR}\, b.$$

The same way as in the proof of Proposition 10.11 we see that μ is an acyclic matching. This matching is complete on the set of simplices in the open star of σ, which of course are precisely the simplices of the difference $\mathcal{K} \setminus \mathrm{dl}_{\mathcal{K}}(\sigma)$. The result then follows from Theorem 10.9. □

Let us now slightly upgrade our argument to deal with the general case.

Proposition 10.13. *Assume $\mathcal{K}$ is an abstract simplicial complex, and assume σ is a simplex of $\mathcal{K}$, such that* $\mathrm{lk}_{\mathcal{K}}(\sigma)$ *is collapsible. Then, again, we have* $\mathcal{K} \searrow \mathrm{dl}_{\mathcal{K}}(\sigma)$.

Proof. By Theorem 10.9 collapsibility of the complex $\mathrm{lk}_{\mathcal{K}}(\sigma)$ implies that there exists a complete acyclic matching on the set of the simplices of $\mathrm{lk}_{\mathcal{K}}(\sigma)$. Let ν denote this matching. For arbitrary simplex $\tau \in \mathcal{K}$ which lies in the open star of σ, we set

$$\mu(\tau) := \nu(\tau \setminus \sigma) \cup \sigma.$$

Assume this matching is not acyclic, and pick a cycle

$$b_1 \succ \mu(b_1) \prec b_2 \succ \mu(b_2) \prec \cdots \prec b_n \succ \mu(b_n) \prec b_1, \text{ where } n \geqslant 2. \tag{10.3}$$

The simplices $b_1, \ldots, b_n$ belong to the open star of σ, which means that $\sigma \subseteq b_i$, for all $i = 1, \ldots, n$. Write $b_1 = a_1 \cup \sigma, \ldots, b_n = a_n \cup \sigma$. Then the cycle in Equation (10.3) becomes

$$a_1 \cup \sigma \succ \nu(a_1) \cup \sigma \prec a_2 \cup \sigma \succ \nu(a_2) \cup \sigma \prec \cdots \prec a_n \cup \sigma \succ \nu(a_n) \cup \sigma \prec a_1 \cup \sigma,$$

and deleting σ everywhere we obtain

$$a_1 \succ \nu(a_1) \prec a_2 \succ \nu(a_2) \prec \cdots \prec a_n \succ \nu(a_n) \prec a_1. \tag{10.4}$$

The existence of the cycle shown by Equation (10.4) contradicts our assumption that ν is an acyclic matching. Thus, assuming that μ was not acyclic was wrong. The statement of the proposition then again follows from Theorem 10.9. □

10.5. Standard subdivisions of collapsible complexes

10.5.1. Stellar subdivision of a collapsible simplicial complex. Next, let us consider a slightly more complicated example. We have seen that the stellar subdivision of a simplex is collapsible. In fact, the following generalization of this fact is true.

Theorem 10.14. *Assume $\mathcal{K}$ is a collapsible abstract simplicial complex, and let α be an arbitrary simplex of $\mathcal{K}$. Then, the abstract simplicial complex $Sd_{\mathcal{K}}(\alpha)$ is collapsible as well.*

The proof of Theorem 10.14 depends on the following lemma.

Lemma 10.15. *Assume n is a positive integer, and $\mathcal{K}$ is an n-simplex, with vertices indexed by the set $[n]$, and let $\mathcal{L}$ be the subcomplex of $\mathcal{K}$ consisting of all simplices except for $[n]$ and $[n-1]$. Let σ be an arbitrary simplex of $\mathcal{K}$. Then we have*

$$\begin{cases} \mathrm{Sd}_{\mathcal{K}}(\sigma) \searrow \mathrm{Sd}_{\mathcal{L}}(\sigma), & \textit{if } \sigma \neq [n-1], [n]; \\ \mathrm{Sd}_{\mathcal{K}}(\sigma) \searrow \mathcal{L}, & \textit{otherwise.} \end{cases} \tag{10.5}$$

Proof. We shall only show the first line of (10.5), leaving the cases $\sigma = [n-1]$ and $\sigma = [n]$ for the exercises.

The simplicial structure of $\mathrm{Sd}_{\mathcal{K}}(\sigma)$ was described in the discussion following Definition 10.4. As was done there, let a denote the barycenter of σ, so the set of the simplices of $\mathrm{Sd}_{\mathcal{K}}(\sigma)$ is the union of the sets $\{\langle a, \tau\rangle \mid \tau \not\supseteq \sigma\}$, and $\{\tau \mid \tau \not\supseteq \sigma\}$.

We need to distinguish two cases. Assume first that $\sigma \subset [n-1]$. Then, we have

$$\mathrm{Sd}_{\mathcal{K}}(\sigma) \setminus \mathrm{Sd}_{\mathcal{L}}(\sigma) = \{\langle a, \tau\rangle \mid \tau \supseteq [n-1] \setminus \sigma,\ \tau \not\supseteq \sigma\}.$$

This is precisely the open star of the simplex $\langle a, [n-1] \setminus \sigma\rangle$. The link of this simplex is a cone with apex at n. It is collapsible, using the matching given by the operation XOR n. It can then be checked directly that his gives the desired collapsing sequence $\mathrm{Sd}_{\mathcal{K}}(\sigma) \searrow \mathrm{Sd}_{\mathcal{L}}(\sigma)$. Alternatively, we can apply Proposition 10.13.

Assume now that n is a vertex of σ, but $\sigma \neq [n]$. In this case, we have

$$\mathrm{Sd}_{\mathcal{K}}(\sigma) \setminus \mathrm{Sd}_{\mathcal{L}}(\sigma) = \{\langle a, \tau\rangle \mid \tau \supseteq [n] \setminus \sigma,\ \tau \not\supseteq \sigma\} \cup \{[n-1]\}.$$

We now start by collapsing the pair $(\langle a, [n-1]\rangle, [n-1])$. This is a legal elementary collapse, as no other simplex of $\mathrm{Sd}_{\mathcal{K}}(\sigma)$ can contain $[n-1]$. After this we just need to collapse the open star of the simplex $\langle a, [n] \setminus \sigma\rangle$. Again, the link of that simplex is a cone with apex at n (it is actually isomorphic to the boundary of a $(\dim \sigma)$-simplex with one top-dimensional simplex removed), so we apply Proposition 10.13 as we did in the previous case. □

Proof of Theorem 10.14. In order to obtain a collapsing sequence for the simplicial complex $\mathrm{Sd}_{\mathcal{K}}(\alpha)$, we start with the collapsing sequence for $\mathcal{K}$, taking it as a blueprint. Then, every time we would have a collapse (σ, τ), in $\mathcal{K}$, such that $\alpha \subseteq \sigma$, we simply replace it with the collapsing sequence either for $\mathrm{Sd}_{\mathcal{K}}(\alpha) \searrow \mathrm{Sd}_{\mathcal{L}}(\sigma)$ or $\mathrm{Sd}_{\mathcal{K}}(\alpha) \searrow \mathcal{L}$, whichever is appropriate, where $\mathcal{L} = \mathcal{K} \setminus \{\tau, \sigma\}$. The existence of the latter collapsing sequence is guaranteed by Lemma 10.15. □

10.5.2. Barycentric subdivision of a collapsible simplicial complex. In general, it is interesting to know which constructions preserve collapsibility. The next theorem tells us that taking barycentric subdivision is one of these constructions.

Theorem 10.16. *Assume $\mathcal{K}$ is a collapsible abstract simplicial complex. Then, the abstract simplicial complex* $\mathrm{Bd}\, \mathcal{K}$ *is collapsible as well.*

More generally, if $\mathcal{K}$ is an abstract simplicial complex and $\mathcal{L}$ its subcomplex such that $\mathcal{K} \searrow \mathcal{L}$, then $\mathrm{Bd}\, \mathcal{K} \searrow \mathrm{Bd}\, \mathcal{L}$.

Corollary 10.17. *The iterated barycentric subdivision of a simplex* $\mathrm{Bd}^t(\Delta^n)$ *is collapsible.*

Before we proceed with the proof of Theorem 10.16 we need the following notion.

Definition 10.18. Assume $\mathcal{K}$ is an abstract simplicial complex. For a simplex $\alpha \in \operatorname{Bd} \mathcal{K}$, $\alpha = \{\sigma_1, \dots, \sigma_t\}$, such that $\sigma_1 \subset \dots \subset \sigma_t$, we shall use the notation $\alpha = (\sigma_1 \subset \dots \subset \sigma_t)$. Furthermore, we call $\sigma_t \in \mathcal{K}$ the *support* of α, and write $\operatorname{supp} \alpha = \sigma_t$.

The following lemma is the analog of Lemma 10.15; it plays the crucial role in the proof of Theorem 10.16.

Lemma 10.19. *Assume* n *is a positive integer, and* $\mathcal{K}$ *is an* n*-simplex, with vertices indexed by the set* $[n]$*. Let* $\mathcal{L}$ *be the subcomplex of* $\mathcal{K}$ *consisting of all simplices except for* $[n]$ *and* $[n-1]$*. Then we have* $\operatorname{Bd} \mathcal{K} \searrow \operatorname{Bd} \mathcal{L}$*.*

Proof. Note that $\operatorname{Bd} \mathcal{L}$ is the subcomplex of $\operatorname{Bd} \mathcal{K}$ consisting of all simplices except for those whose support is $[n]$ or $[n-1]$. Let C denote the set of simplices $\operatorname{Bd} \mathcal{K} \setminus \operatorname{Bd} \mathcal{L}$. We define the matching $\mu : C \to C$ as follows. The first partial rule says

$$\mu(\sigma_1 \subset \dots \subset \sigma_k) := \begin{cases} \sigma_1 \subset \dots \subset \sigma_k \subset [n], & \text{if } \sigma_k = [n-1]; \\ \sigma_1 \subset \dots \subset \sigma_{k-1}, & \text{if } \sigma_{k-1} = [n-1]. \end{cases} \tag{10.6}$$

For the last line note that when $\sigma_{k-1} = [n-1]$, we automatically have $\sigma_k = [n]$. This will match completely all simplices α in C which have $\operatorname{supp} \alpha = [n-1]$ with those which have $\operatorname{supp} \alpha = [n]$ and $\sigma_{k-1} = [n-1]$.

What remains are the simplices $\alpha = (\sigma_1 \subset \dots \subset \sigma_k)$ for which $\sigma_k = [n]$ and $\sigma_{k-1} \neq [n-1]$. Assume α is such a simplex, and let $h(\alpha)$ be the maximal index $1 \leqslant h(\alpha) \leqslant k$ such that $n \notin \sigma_{h(\alpha)}$. If $n \in \sigma_i$ for all $1 \leqslant i \leqslant k$, then we set $h(\alpha) := 0$ and use $\sigma_0 = \emptyset$ as a default value.

The next rule completes our definition of μ:

$$\mu(\sigma_1 \subset \dots \subset \sigma_k)$$
$$:= \begin{cases} \sigma_1 \subset \dots \subset \sigma_{h(\alpha)} \subset \sigma_{h(\alpha)} \cup n \subset \dots \subset \sigma_k, & \text{if } |\sigma_{h(\alpha)+1}| \geqslant |\sigma_{h(\alpha)}| + 2; \\ \sigma_1 \subset \dots \subset \sigma_{h(\alpha)} \subset \sigma_{h(\alpha)+2} \subset \dots \subset \sigma_k, & \text{if } |\sigma_{h(\alpha)+1}| = |\sigma_{h(\alpha)}| + 1. \end{cases}$$

It is clear that μ is a complete matching on the set C. Before proceeding, we make the following observations:

- by definition of μ, we have $\operatorname{supp} b = [n]$ for all $b \in M^{\uparrow}(C)$;
- for all $b \in C$, we have $h(\mu(b)) = h(b)$, since μ adds or deletes a set containing n;
- if $b \prec c$ in $\operatorname{Bd} \mathcal{K}$, then $h(b) \leqslant h(c)$.

Let us now show that μ is acyclic. Assume it is not, and take a cycle

$$b_1 \succ \mu(b_1) \prec b_2 \succ \mu(b_2) \prec \cdots \prec b_t \succ \mu(b_t) \prec b_1,$$

where $n \geqslant 2$, and all the simplices are distinct. By what is said above, we have

$$h(b_1) = h(\mu(b_1)) \leqslant h(b_2) = h(\mu(b_2)) \leqslant \cdots \leqslant h(b_t) = h(\mu(b_t)) \leqslant h(b_1),$$

which of course implies

$$h(b_1) = h(\mu(b_1)) = h(b_2) = h(\mu(b_2)) = \cdots = h(b_t) = h(\mu(b_t)) = h(b_1).$$

To start with, assume that all of the matched pairs $(b_1, \mu(b_1)), \dots, (b_t, \mu(b_t))$ are of the first type, as defined in Equation (10.6); in other words, $\operatorname{supp}\mu(b_1) = \cdots = \operatorname{supp}\mu(b_t) = [n-1]$. Let us say $\mu(b_1) = (\sigma_1 \subset \cdots \subset \sigma_{d-1} \subset [n-1])$. Since $\operatorname{supp} b_1 = \cdots = \operatorname{supp} b_t = [n]$, we get

$$b_1 = (\sigma_1 \subset \cdots \subset \sigma_{d-1} \subset [n-1] \subset [n]) = b_2,$$

which obviously is a contradiction.

We can therefore assume without loss of generality that $\operatorname{supp}\mu(b_1) \neq [n-1]$, which of course implies $\operatorname{supp}\mu(b_1) = [n]$. Since $\mu(b_1) \in M^{\downarrow}(C)$, assuming $\mu(b_1) = (\sigma_1 \subset \cdots \subset \sigma_d)$, we must have

$$b_1 = (\sigma_1 \subset \cdots \subset \sigma_h \subset \sigma_h \cup n \subset \sigma_{h+1} \subset \cdots \subset \sigma_d),$$

where $h = h(\mu(b_1))$. We furthermore have assumed that $b_1 \succ \mu(b_t)$ and derived $h(b_1) = h(\mu(b_t))$. This forces us to conclude that

$$\mu(b_t) = (\sigma_1 \subset \cdots \subset \sigma_h \subset \sigma_h \cup n \subset \sigma_{h+1} \subset \cdots \subset \sigma_{j-1} \subset \sigma_{j+1} \subset \cdots \subset \sigma_d),$$

for some $h + 1 \leqslant j \leqslant d$. However, this means that $\mu(b_t) \in M^{\uparrow}(C)$, again yielding a contradiction.

All in all, we can conclude that the matching μ is acyclic. The result now follows from Theorem 10.9. $\square$

Proof of Theorem 10.16. We shall show the second statement. Assume $\mathcal{K}$ is an abstract simplicial complex and $\mathcal{L}$ its subcomplex such that $\mathcal{K} \searrow \mathcal{L}$. Take some sequence of elementary collapses $(\sigma_1, \tau_1), \dots, (\sigma_t, \tau_t)$ leading from $\mathcal{K}$ to $\mathcal{L}$. During the elementary collapse (σ_k, τ_k) for some $1 \leqslant k \leqslant t$ we remove two simplices: σ_k and τ_k. Passing on to the barycentric subdivisions, we would like to come from $\operatorname{Bd}\mathcal{K}$ to $\operatorname{Bd}\mathcal{L}$ using essentially the same collapsing steps. The difference is that now instead of removing σ_k and τ_k we would like to delete the whole set of simplices $\operatorname{Bd}\sigma_k \cup \operatorname{Bd}\tau_k$. Here $\operatorname{Bd}\sigma_k$ denotes the set of simplices of $\operatorname{Bd}\mathcal{K}$ whose support is equal to σ_k, and similarly for $\operatorname{Bd}\tau_k$. Fortunately, this is precisely the statement of Lemma 10.19, so we are done. $\square$

10.6. Standard chromatic subdivision

10.6.1. Standard chromatic subdivision of a simplex. Let n be a natural number, and again let Δ^n be the standard n-simplex. The abstract simplicial complex $\chi(\Delta^n)$ is a pure n-dimensional abstract simplicial complex defined as follows:

- the vertices of $\chi(\Delta^n)$ are indexed by all pairs (p, V), such that $V \subseteq [n]$, and $p \in V$;
- the n-dimensional simplices of $\chi(\mathcal{K})$ are formed by all sets of vertices $\{(0, V_0), (1, V_1), \dots, (n, V_n)\}$ satisfying the following axioms:
  (i) for all $i, j \in [n]$, we have either $V_i \subseteq V_j$ or $V_j \subseteq V_i$;
  (ii) for all $i, j \in [n]$, if $i \in V_j$, then $V_i \subseteq V_j$.

In particular, note that $\chi(\Delta^n)$ has $2^n(n+1)$ vertices in total.

Definition 10.20. The abstract simplicial complex $\chi(\Delta^n)$ is called the *standard chromatic subdivision* of Δ^n.

The standard chromatic subdivision of a simplicial complex is similar to the barycentric subdivision. Yet, there is a crucial difference, which also explains the etymology of the term. Imagine that we are interested in coloring the vertices of a simplicial complex $\mathcal{K}$, so that if two vertices are connected by an edge, they must get different colors. Clearly, we will need at least $\dim \mathcal{K} + 1$ colors, since all the vertices of any simplex of maximal dimension will need to be covered differently. However, if we did succeed to color $\mathcal{K}$ with $\dim \mathcal{K} + 1$ colors, taking the barycentric subdivision will ruin it, since the color assigned to the barycenter of any top-dimensional simplex σ cannot be any of the old ones, as this barycenter is connected to all the vertices of σ. This is not the case for the standard chromatic subdivision. On the contrary, given a valid coloring of the vertices of $\mathcal{K}$, the vertices of the simplicial complex $\chi(\mathcal{K})$ can be colored using the same of colors.

It is useful to obtain an alternative combinatorial description of the entire simplicial structure of $\chi(\Delta^n)$, including the boundary operator.

Definition 10.21. A *partial ordered set partition* of the set $[n]$ is a pair of ordered set partitions of non-empty subsets of $[n]$, $\sigma = ((A_1, \dots, A_t), (B_1, \dots, B_t))$, which have the same number of parts and are subject to the following additional conditions:

- for all $1 \leqslant i \leqslant t$, we have $B_i \subseteq A_i$;
- the sets A_i are disjoint.

Given such a partial ordered set partition σ, we introduce the following terminology.

- The union $A_1 \cup \cdots \cup A_t$ is called the *carrier set* of σ, and is denoted by $R(\sigma)$.
- The union $B_1 \cup \cdots \cup B_t$ is called the *color set* of σ, and is denoted by $C(\sigma)$.
- The *dimension* of σ is defined to be $|C(\sigma)|-1$, and is denoted $\dim \sigma$.

When appropriate, we shall also write

$$\sigma = \begin{array}{|c|c|c|}\hline A_1 & \dots & A_t \\ \hline B_1 & \dots & B_t \\ \hline\end{array}, \tag{10.7}$$

which we shall call the *table form* of σ.

We note that both nodes (A, x), for $A \subseteq [n]$, as well as ordered set partitions of $[n]$ are special cases of partial ordered set partitions of $[n]$. Indeed, a node (A, x), such that $A \subseteq [n]$, corresponds to the somewhat degenerate partial ordered set partition of $[n]$,

$$\sigma = \begin{array}{|c|}\hline A \\ \hline x \\ \hline\end{array},$$

whereas an ordered set partition $(A_1, \dots, A_t)$ corresponds to the partial ordered set partition of $[n]$,

$$\sigma = \begin{array}{|c|c|c|}\hline A_1 & \dots & A_t \\ \hline A_1 & \dots & A_t \\ \hline\end{array},$$

i.e., any partial ordered set partition $((A_1, \dots A_t), (B_1, \dots, B_t))$, such that $A_i = B_i$ for all i, and $A_1 \cup \cdots \cup A_t = [n]$.

Each partial ordered set partition has a non-empty color set, which in turn is contained in its carrier set. The nodes correspond to the partial ordered set partitions with minimal color set, consisting of just one element, and ordered set partitions correspond to the partial ordered set partitions with maximal color set, namely the whole set $[n]$.

Definition 10.22. Assume we are given a partial ordered set partition of the set $[n]$, say $\sigma = ((A_1, \dots, A_t), (B_1, \dots, B_t))$, such that $\dim \sigma \geqslant 1$, and we are also given an element $x \in C(\sigma)$, say $x \in B_k$, for some $1 \leqslant k \leqslant t$. To define the *deletion of* x *from* σ we consider three different cases.

Case 1. If $|B_k| \geqslant 2$, then the deletion of x from σ is set to be

$$((A_1, \dots, A_t), (B_1, \dots, B_{k-1}, B_k \setminus x, B_{k+1}, \dots, B_t)).$$

Case 2. If $|B_k| = 1$ and $k \leqslant t-1$, then the deletion of x from σ is set to be

$$((A_1, \dots, A_{k-1}, A_k \cup A_{k+1}, \dots, A_t), (B_1, \dots, B_{k-1}, B_{k+1}, \dots, B_t)).$$

Case 3. If $|B_k| = 1$ and $k = t$, then the deletion of x from σ is set to be

$$((A_1, \dots, A_{t-1}), (B_1, \dots, B_{t-1})).$$

We denote the deletion of x from σ by $\mathrm{dl}(\sigma, x)$.

The deletion of an element corresponds to the boundary relation in the standard chromatic subdivision. (See Figure 10.7.)

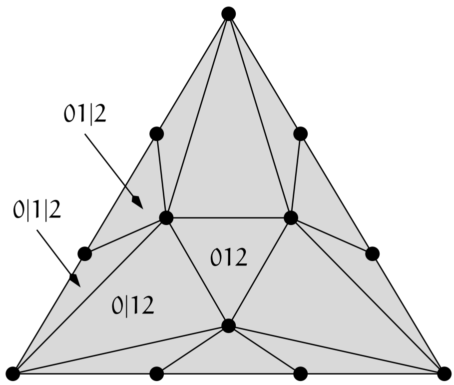

Figure 10.7. Standard chromatic subdivision of a triangle.

Theorem 10.23. *The standard chromatic subdivision of a simplex is collapsible.*

Proof. Consider an n-simplex $\mathcal{K}$, whose set of vertices is indexed by $[n]$. Let B be the set of all simplices

$$\sigma = \begin{array}{|c|c|c|} \hline A_1 & \dots & A_t \\ \hline A_1 & \dots & A_t \\ \hline \end{array},$$

such that $C(\sigma) = A_1 \cup \cdots \cup A_t \neq [n]$. For each $\sigma \in B$, set

$$\lambda(\sigma) := \begin{array}{|c|c|c|c|} \hline A_1 & \dots & A_t & [n] \setminus C(\sigma) \\ \hline A_1 & \dots & A_t & [n] \setminus C(\sigma) \\ \hline \end{array}.$$

It is then easy to see that the face poset of $\chi(\mathcal{K})$ decomposes as a disjoint union of intervals $[\sigma, \lambda(\sigma)]$, where $\sigma \in B$. Each such interval corresponds to a (possibly non-elementary) collapse. The collapses can be performed in any order, so that the cardinality of $C(\sigma)$ does not decrease. $\square$

10.6.2. Standard chromatic subdivision of a collapsible simplicial complex.

Definition 10.24. Let $\mathcal{K}$ be an arbitrary abstract simplicial complex. A new simplicial complex $\chi(\mathcal{K})$, called the *standard chromatic subdivision* of $\mathcal{K}$, is defined as follows:

- the simplices of $\chi(\mathcal{K})$ are indexed by all partial ordered set partitions
$$\sigma = \begin{array}{|c|c|c|}\hline A_1 & \dots & A_t \\ \hline B_1 & \dots & B_t \\ \hline\end{array},$$
such that $A_1 \cup \dots \cup A_t$ is a simplex of $\mathcal{K}$;
- the boundary relation is encoded by the deletion of elements, as described in Definition 10.22.

Theorem 10.25. *If $\mathcal{K}$ is an abstract simplicial complex and $\mathcal{L}$ its subcomplex such that $\mathcal{K} \searrow \mathcal{L}$, then $\chi(\mathcal{K}) \searrow \chi(\mathcal{L})$. In particular, the standard chromatic subdivision of a collapsible simplicial complex is collapsible.*

In line with our previous arguments, the collapsing sequence required in Theorem 10.25 can be obtained by an iterated application of the following lemma.

Lemma 10.26. *Let $\mathcal{K}$ denote the n-simplex, with the vertex set $[n]$, and let Λ denote the subcomplex of $\mathcal{K}$ obtained by an elementary collapse $([n], [n-1])$. Then, we have $\chi(\mathcal{K}) \searrow \chi(\Lambda)$.*

Proof. The proof proceeds in three stages. First, we perform all the collapses
$$\left(\begin{array}{|c|c|c|c|}\hline A_1 & \dots & A_t & [n] \setminus (A_1 \cup \dots \cup A_t) \\ \hline B_1 & \dots & B_t & n \\ \hline\end{array}, \begin{array}{|c|c|c|}\hline A_1 & \dots & A_t \\ \hline B_1 & \dots & B_t \\ \hline\end{array}\right),$$
where $n \notin A_1 \cup \dots \cup A_t$.

Second, collapse the open star of the vertex $([n], n)$, using Propoition 10.13.

Finally, arrange the remaining simplices $([n], S)$, where $S \subseteq [n-1]$, in such an order that the cardinality of S does not increase, and use Proposition 10.13 again to delete the open stars of these simplices in this chosen order. Details are left as an exercise. □

Proof. Given a collapsing sequence leading from $\mathcal{K}$ to $\mathcal{L}$, we simply replace each elementary collapse in $\mathcal{K}$ by the corresponding sequence of elementary collapses in $\chi(\mathcal{K})$, whose existence is guaranteed by Lemma 10.26. □

Corollary 10.27. *The iterated standard chromatic subdivision of a simplex $\chi^t(\Delta^n)$ is collapsible.*

Proof. By Theorem 10.23 the standard chromatic subdivision of a simplex is collapsible. Now, repeated application of the second statement of Theorem 10.25 yields the result. □

10.7. Combinatorial collapsing sequences

10.7.1. Closure operators in posets. Let us now describe a framework which is classical in combinatorial topology, and which yields a sequence of simplicial collapses.

Definition 10.28. Let P be a partially ordered set. A *closure operator* on P is a poset map $\varphi : P \to P$ satisfying the following conditions:

(1) $\varphi(x) \geqslant x$, for all $x \in P$;
(2) $\varphi^2 = \varphi$.

Assume $\varphi : P \to P$ is a closure operator, and let $\varphi(P)$ denote its image $\{\varphi(x) \,|\, x \in P\}$. This is of course a partially ordered set as well, with the partial order induced by that of P. Furthermore, by condition (2) in Definition 10.28, the restriction of φ to the poset $\varphi(P)$ is simply the identity map. Moreover, we obtain the following implication:

$$\begin{cases} x \in P \\ y \in \varphi(P) \\ x \leqslant y \end{cases} \implies \quad \varphi(x) \leqslant y. \tag{10.8}$$

Indeed, since φ is a poset map, the inequality $x \leqslant y$ implies the inequality $\varphi(x) \leqslant \varphi(y)$. On the other hand, $\varphi(y) = y$, since $y \in \varphi(P)$, so we obtain $\varphi(x) \leqslant y$.

The following theorem provides a connection between closure operators and simplicial collapsing sequences.

Theorem 10.29. *Let P be a finite partially ordered set, and let $\varphi : P \to P$ be a closure operator. Then there exists a collapsing sequence from the order complex $\Delta(P)$ to the order complex $\Delta(\varphi(P))$.*

Proof. Clearly, the simplicial complex $\Delta(\varphi(P))$ is a subcomplex of $\Delta(P)$. Let us define an acyclic matching on the simplices of $\Delta(P)$ such that the set of critical simplices is precisely $\Delta(\varphi(P))$.

Recall that simplices of $\Delta(P)$ are chains $c = (x_1 < \dots < x_t)$. Assume that c is not a chain in $\varphi(P)$, and let k be the largest index between 1 and t, such that $x_k \notin \varphi(P)$. Note that if $k \leqslant t-1$, then $x_{k+1} \in \varphi(P)$, so, by Equation (10.8), we have $x_k < \varphi(x_k) \leqslant \varphi(x_{k+1})$.

We now define a matching by setting

$$\mu(c) := c \,\mathrm{XOR}\, \varphi(x_k),$$

where XOR is the *exclusive OR* operation defined in Definition 10.10.

Let us see that μ is an acyclic matching. Consider a cycle

$$b_1 \prec \mu(b_1) \succ b_2 \prec \mu(b_2) \succ \dots \succ b_n \prec \mu(b_n) \succ b_1.$$

For a chain $c = (x_1 < \cdots < x_t)$ in P, let $h(b)$ denote the number of x_i's which belong to the image $\varphi(P)$. On one hand, we have $h(\mu(b_i)) = h(b_i) + 1$, for $i = 1, \dots, n$, since when the matching operation adds an element, it adds an element from the image of φ. On the other hand, $h(\mu(b_i)) - 1 \leqslant h(b_{i+1})$, for $i = 1, \dots, n$,[4] simply because when passing from $\mu(b_i)$ to b_{i+1}, we may delete at most one element from the image of φ. So, in total, we have

$$h(b_1) = h(\mu(b_1)) - 1 \leqslant h(b_2) = \cdots \leqslant h(b_n) = h(\mu(b_n)) - 1 \leqslant h(b_1).$$

We then have $h(b_1) = \cdots = h(b_n)$. Assume $\mu(b_1) = b_1 \cup \varphi(x_k)$. When passing from $\mu(b_1)$ to b_2 we must delete an element from the image of φ, so we cannot delete x_k. But then x_k is still the largest element of b_2 which lies outside of the image of φ, so $\mu(b_2) = b_2 \cup \varphi(x_k)$. This, of course, is impossible, since $b_1 \neq b_2$, so $\varphi(x_k) \in b_2$.

We conclude that μ is acyclic. Clearly, the set of those simplices which are not matched by μ coincides with the set of simplices of $\Delta(\varphi(P))$, so the proposition follows from Theorem 10.9. $\square$

10.7.2. Order complexes of posets with a join-transversal.

Definition 10.30. An element α of P is called a *join-transversal* if for any other element $x \in P$ the join of x and α exists.

There is an easy way to produce a poset with a join-transversal. Start with any poset P and an element $\alpha \in P$, and then simply delete all $x \in P$ such that the join of α and x does not exist. It is an easy exercise to see that α will be a join-transversal in the resulting poset.

Theorem 10.31. *Assume P is a poset possessing a join-transversal; then the order complex of P is collapsible.*

Proof. This is a simple corollary of Theorem 10.29. Specifically, assume $\alpha \in P$ is a join-transversal. We define $\varphi : P \to P$ by setting $\varphi(x) := x \vee \alpha$. Obviously it is a poset map and $\varphi(x) \geqslant x$. Also $(x \vee \alpha) \vee \alpha = x \vee \alpha$, so $\varphi(\varphi(\alpha)) = \varphi(\alpha)$, making φ a closure operator. The image $\varphi(P)$ is precisely $P_{\geqslant \alpha}$, since, on one hand $x \vee \alpha \geqslant \alpha$, for all $x \in P$, and on the other hand, $x \vee \alpha = x$, whenever $x \geqslant \alpha$, so each element of $P_{\geqslant \alpha}$ lies in $\varphi(P)$.

By Theorem 10.29 this means that there exists a collapsing sequence from $\Delta(P)$ to $\Delta(P_{\geqslant \alpha})$. However, the simplicial complex $\Delta(P_{\geqslant \alpha})$ is a cone with apex at α, which is collapsible. Therefore $\Delta(P)$ is collapsible as well. $\square$

---

[4]As usual, we use the convention $b_{n+1} = b_1$.

10.7.3. Removing dominating vertices in independence complexes.

Recall that a set of vertices of a graph is called *independent* if no two of them are connected by an edge. Now, given a graph G, its *independence complex* $\mathrm{Ind}(G)$ has vertices of G as vertices, and independent sets of G as simplices.

Definition 10.32. Assume we are given a graph G and two vertices of G, say v and w. We say that w *dominates* v if any vertex which is adjacent to v is also adjacent to w.[5]

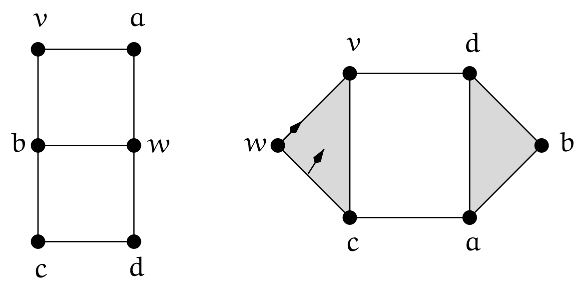

Figure 10.8. On the left-hand side we see a graph G with vertex w dominating the vertex v. On the right-hand side we see the independence complex $\mathrm{Ind}(G)$ and the collapsing sequence used in the proof of Proposition 10.33.

Proposition 10.33. *Assume we are given a graph G and two vertices v and w such that w dominates v; then there is a sequence of simplicial collapses reducing $\mathrm{Ind}(G)$ to $\mathrm{Ind}(G \setminus w)$.*

Proof. The following observation is crucial for defining the collapsing sequence: *if A is an independent set in G, and $w \in A$, then also $A \cup v$ is an independent set in G.* Indeed, if $A \cup v$ is not an independent set, then there exists $x \in A$, such that the edge (x, v) belongs to G. Since w dominates v, also the edge (x, w) must be an edge of G, contradicting the fact that A is an independent set.

Equipped with this fact, we simply define the matching $\mu(A) := A \,\mathrm{XOR}\, v$, for all $A \in \mathrm{Ind}(G)$, such that $w \in A$. This is well-defined and obviously acyclic. Figure 10.8 illustrates our argument. $\square$

Proposition 10.33 can be effectively used to reduce independence complexes of graphs to independence complexes of potentially much smaller graphs. For example, when G is a path of n vertices, as in Exercise (2), then Proposition 10.33 allows us to obtain very extensive information on $\mathrm{Ind}(G)$.

[5] When the vertices are persons, and edges record who knows whom, w dominates v if it knows everybody whom v knows.

10.7.4. Complexes of disconnected graphs and the order complex of the partition lattice. There is a standard way to associate a family of simplicial complexes to any graph property Λ which is closed under the deletion of edges. In other words, if G has property Λ and e is an arbitrary edge of G, then also $G \setminus e$ will have the property Λ.

Definition 10.34. Let Λ be a graph property closed under the deletion of edges, and let n be an arbitrary integer, $n \geqslant 2$. We define a simplicial complex $\mathrm{GP}_n(\Lambda)$ as follows:

(1) The vertices of $\mathrm{GP}_n(\Lambda)$ are indexed by all ordered pairs (i, j), where $1 \leqslant i < j \leqslant n$; there are $\binom{n}{2}$ vertices.
(2) Given a set σ of vertices of $\mathrm{GP}_n(\Lambda)$, we can construct a graph G_σ, whose vertices are indexed $1, \dots, n$, and edges are listed in σ. We now say that σ is a simplex of $\mathrm{GP}_n(\Lambda)$ if and only if the graph G_σ has the property Λ.

Varying n, we get an infinite family of abstract simplicial complexes.

Colloquially we shall refer to the simplicial complexes $\mathrm{GP}_n(\Lambda)$ as *complexes of all graphs with the property* Λ. In principle, one can take any graph property which is closed under the deletion of edges, such as planarity, or colorability with a certain number of colors. In general however, the resulting simplicial complexes can be hard to analyze.

Remark 10.35.

(1) Definition 10.34 can be easily modified to describe abstract simplicial complexes associated to those graph properties of *directed graphs*, which are closed under edge deletion.
(2) Another possible generalization of Definition 10.34 is to consider *hypergraphs* instead of graphs.

Here, we consider the simplicial complexes of disconnected graphs. These are of course well-defined, since a deletion of an edge from a disconnected graph again yields a disconnected graph. For brevity we shall let Disc_n denote the abstract simplicial complex of all disconnected graphs on n vertices. For $n = 2$, this simplicial complex is empty. For $n = 3$ it consists of 3 isolated vertices. The example $n = 4$ is shown in Figure 10.9. This simplicial complex can be visualized as taking every other triangle in an octahedron and then connecting the 3 pairs of opposite vertices by edges. It has the homotopy type of a wedge of 6 circles.

The case $n = 5$ is slightly more difficult. The simplicial complex Disc_5 has dimension 5. It has 5 maximal simplices of dimension 5 and 10 maximal simplices of dimension 3. With a little bit of extra effort one can figure out that Disc_5 is homotopy equivalent to a wedge of 24 spheres of dimension 2.

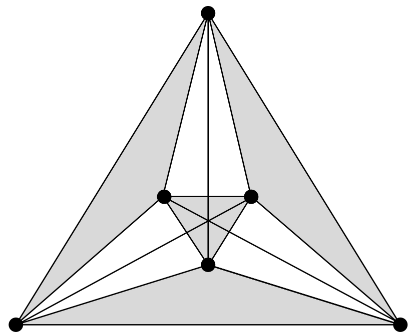

Figure 10.9. The complex of disconnected graphs on 4 vertices.

The observations made for the lower values of n carry over to the higher-dimensional case. The maximal simplices of Disc_n are indexed by partitions of the set $\{1, \ldots, n\}$ into two non-empty subsets. Accordingly, when these parts have cardinalities k and $n-k$, the corresponding simplex has dimension $\binom{k}{2} + \binom{n-k}{2} - 1$, so the dimension of Disc_n is equal to $\binom{n-1}{2} - 1$.

As a contrast, its homology is concentrated in a much lower dimension. As a matter of fact, it turns out that the simplicial complex Disc_n is homotopy equivalent to the wedge of $(n-1)!$ spheres of dimension $n-3$. We will see this in Chapter 11, however the fact that the homology in higher dimensions is trivial can be shown using collapsing sequences. In order to do this, we need another object, frequently met in combinatorial topology, which we now proceed to introduce.

Definition 10.36. Let n be an integer, $n \geqslant 3$. The *partition lattice* Π_n is the poset whose elements are all set partitions of the set $\{1, \ldots, n\}$, with the partial order given by refinement. (See Figure 10.10.)

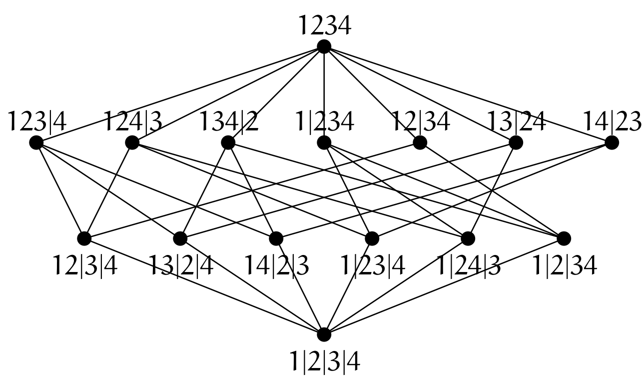

Figure 10.10. The partition lattice Π_4.

The idea behind Definition 10.36 is to encode the combinatorics of coincidences. For example, the Euclidean space $\mathbb{R}^n$ consists of all n-tuples of real numbers. We can stratify this space by geometric locuses of points with various sets of coordinates coinciding. Geometrically this is known as the *braid arrangement* and the stratification poset is precisely the partition lattice. More generally, the partition lattice describes the combinatorics of the standard stratification of the ordered configuration space of n points in an arbitrary topological space.

The partition lattice Π_n has a maximal element corresponding to the partition consisting of a single block $(1, \dots, n)$, and a minimal element corresponding to the partition $(1) \dots (n)$ consisting of singletons[6] only.

In general, we said that $\hat{0}$ should denote the minimal element in a poset, so it is close at hand to extend this notation, and let $\hat{1}$ stand for the maximal element. In many natural situations, as is the case for the partition lattice, the considered poset will have both the maximal and the minimal elements. When considering the associated complex, it is useful to exclude both, so as to avoid the double cone, which we would get otherwise. This makes the next definition natural.

Definition 10.37. Assume P is a partially ordered set, having both the minimal element $\hat{0}$, as well as the maximal element $\hat{1}$. We call the order complex $\Delta(P \setminus \{\hat{0}, \hat{1}\})$ the *reduced order complex* of P, and denote it by $\widetilde{\Delta}(P)$.

The following proposition provides a relation between the complex of disconnected graphs with the reduced order complex of partition lattice.

Proposition 10.38. *For every $n \geqslant 2$, there exists a collapsing sequence leading from the barycentric subdivision of the simplicial complex of all disconnected graphs on n vertices to the reduced order complex of the partition lattice Π_n, i.e., we have*

$$\operatorname{Bd}\operatorname{Disc}_n \searrow \widetilde{\Delta}(\Pi_n).$$

Proof. Let P be the face poset of Disc_n, and let $\varphi : P \to P$ map each graph to its transitive closure. The graphs in the image of φ all consist of disjoint unions of complete graphs, which are in a clear one-to-one correspondence with partitions of the vertex set. Since the elements of P encode disconnected graphs, we will never get the partition consisting of a single block, so we stay within $\Pi_n \setminus \{\hat{1}\}$. The statement now follows from Theorem 10.29. Figure 10.11 illustrates our argument. □

[6]Recall that *singletons* is the short term used for sets consisting of one element.

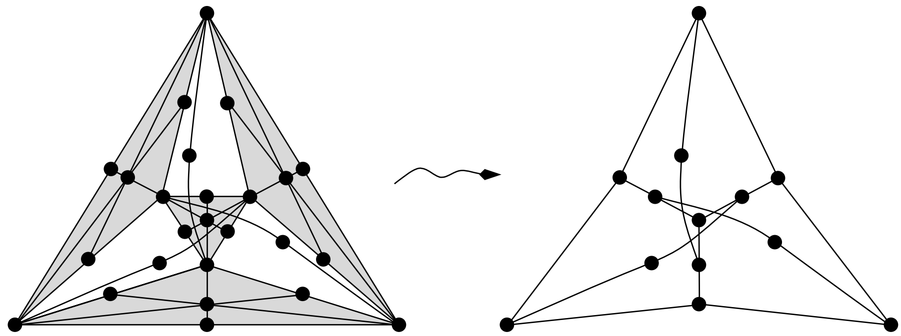

Figure 10.11. Collapsing the barycentric subdivision of the complex of disconnected graphs on 4 vertices to $\Delta(\Pi_4)$.

Exercises

(1) Prove Proposition 10.8.

(2) Compute the homotopy type of independence graphs of paths. Do the same for cycles.

(3) Characterize all graphs such that the iterative removal of dominating vertices gives a disjoint union of isolated vertices or edges. Use this to compute the homotopy type of the corresponding independence complexes.

(4) A finite abstract simplicial complex $\mathcal{K}$ is called *nonevasive* if it is either a point or contains a vertex v such that both the deletion of v as well as the link of v are non-evasive.
   (a) Show that the reduced Euler characteristic of a non-evasive simplicial complex is equal to 0.
   (b) Show that a non-evasive simplicial complex is collapsible.
   (c) Can you find a collapsible simplicial complex which is evasive (i.e., not non-evasive)?

(5) Assume $\mathcal{K}$ and $\mathcal{M}$ are abstract simplicial complexes. Let $d_{\text{coll}}(\mathcal{K},\mathcal{M})$ denote the minimal number of elementary simplicial collapses and expansions which are needed to transform $\mathcal{K}$ to $\mathcal{M}$. We call $d_{\text{coll}}(\mathcal{K},\mathcal{M})$ the *collapsing distance* between $\mathcal{K}$ and $\mathcal{M}$.
   (a) Show that d_{coll} is a metric.
   (b) For any $k \geqslant 3$, let C_k denote the cyclic graph with k edges. Assume $n \geqslant m \geqslant 3$. Show that $d_{\text{coll}}(C_m, C_n) = 3(n-m)$.
   (c) Assume $k, m \geqslant 3$. Let X_k be the *star graph* with k edges: its set of vertices is $[k]$, and its set of edges is $\{(0,v) \mid 1 \leqslant v \leqslant k\}$. Let P_m be the path graph with m edges: its set of vertices is $[m]$, and

its set of edges is $\{(v, v+1) \mid 0 \leqslant v \leqslant m-1\}$. Derive a formula for $d_{\mathrm{coll}}(X_k, P_m)$.

(6) Let a group G act on an abstract simplicial complex $\mathcal{K}$. Assume that removing a pair of simplices (σ, τ) is a simplicial collapse. Assume furthermore that we have the inclusion of stabilizers $\operatorname{stab}\sigma \subseteq \operatorname{stab}\tau$, and consider the set of pairs $S = \{(g\sigma, g\tau) \mid g \in G\}$.

Show that removing each element of S is a simplicial collapse, and that these simplicial collapses are disjoint and can be performed in an arbitrary order, independently of each other.

In a situation like this, we say that removing the entire set S is an *equivariant simplicial collapse*.

(7) Assume $\mathcal{K}$ is a finite abstract simplicial complex, and the group G acts on $\mathcal{K}$. The complex $\mathcal{K}$ is called *equivariantly collapsible* if there exists a sequence of equivariant collapses (defined in Exercise (6)), with respect to this G-action, resulting in a void simplicial complex.
   (a) Assume $n \geqslant 2$. Show that the stellar subdivision of an n-simplex is equivariantly collapsible with respect to the natural vertex permutation action by S_{n+1}.
   (b) Show that the barycentric subdivision of an n-simplex is equivariantly collapsible with respect to the same action.
   (c) Show that the chromatic subdivision of an n-simplex is equivariantly collapsible with respect to the same action.

(8) Show the second line in Equation (10.5).

(9) Complete the proof of Lemma 10.26 by checking the correctness of the described collapsing sequence.

Chapter 12

Explicit Homology Classes Associated to Critical Cells

12.1. The graph $G_d(\mu)$ and its use

Let us now describe how in some cases discrete Morse theory allows us to find explicit homology classes.

Definition 12.1. Assume we are given an abstract simplicial complex $\mathcal{K}$ and an acyclic matching μ. For every d, between 0 and $\dim \mathcal{K}$, the directed graph $G_d(\mu)$ is defined as follows:

- the vertices of $G_d(\mu)$ are indexed by the d-dimensional simplices of $\mathcal{K}$;
- the edges of $G_d(\mu)$ are given by the rule: (α, β) is an edge of $G_d(\mu)$ if and only if $\mu(\beta)$ is defined, and $\alpha \succ \mu(\beta)$; see Figure 12.1.

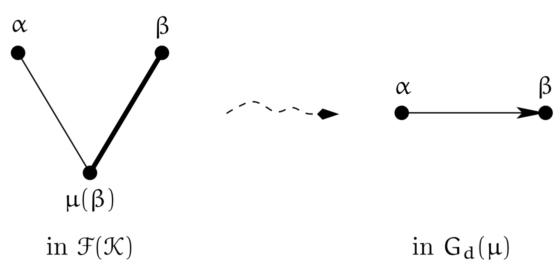

Figure 12.1. The rule defining the edges of the directed graph $G_d(\mu)$.

Obviously, the acyclicity of the matching μ is equivalent to the acyclicity of the directed graph $G_d(\mu)$.

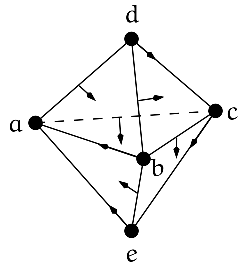

Figure 12.2. A simplicial complex with an acyclic matching.

It is a well-known fact in graph theory that the vertex set of a finite acyclic directed graph $G_d(\mu)$ can be decomposed into layers, that is, represented as a disjoint union $V_0 \cup \cdots \cup V_t$, such that

(1) for any $\alpha \in V_i$, there exists $\beta \in V_{i-1}$, such that (α, β) is an edge of $G_d(\mu)$,

(2) for any $\alpha \in V_i$, $\beta \in V_j$, such that (α, β) is an edge of $G_d(\mu)$, we have $i > j$.

In other words, V_0 consists of all the *sinks*, the vertices with outdegree equal to 0; V_1 consists of all the vertices with edges pointing to sinks only, and, in general, for each i, the set V_i consists of vertices α such that the longest path from α to a sink has length i. Another way to say this is that the sets $V_1, \dots, V_t$ are precisely the sets which would be produced by the breadth-first search algorithm, starting from the set of all sinks V_0, and tracing the edges in the opposite direction.

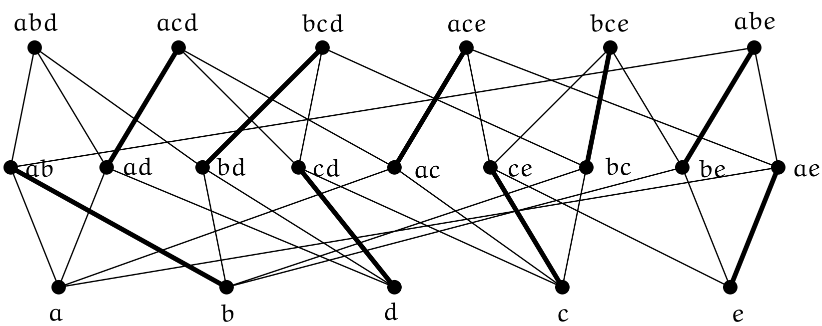

Figure 12.3. The acyclic matching on the face poset (with $\emptyset$ omitted) of the simplicial complex from Figure 12.2.

An instance of how this construction works is shown in Figures 12.2 to 12.4. Note that, in general, edges are allowed to skip levels, by which we mean that we may have edges (α, β), such that $\alpha \in V_i$, $\beta \in V_j$, and $i \geqslant j+2$. For example, in the graph in Figure 12.4 the edge between abd and acd skips a level.

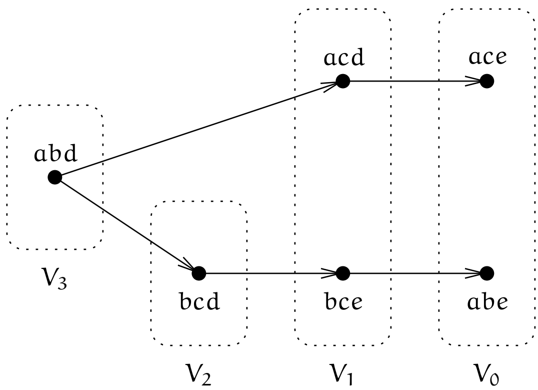

Figure 12.4. The directed graph $G_2(\mu)$, for the acyclic matching μ from Figures 12.2 and 12.3, and its decomposition into layers $V_0, \dots, V_3$.

Let μ be the acyclic matching, and, as previously defined, let R, $M^{\uparrow}$, and $M^{\downarrow}$ be the corresponding decomposition. Let us now consider consequences for the d-chains of $\mathcal{K}$.

Proposition 12.2. *Assume* $\sigma \in C_d(\mathcal{K})$ *and* $\operatorname{supp} \sigma \subseteq M^{\uparrow}$.

(1) *If* σ *is a cycle, then* $\sigma = 0$.

(2) *More generally, if* $\sigma \neq 0$, *then*

$$\operatorname{supp} \partial\sigma \cap \mu(\operatorname{supp} \sigma) \neq \emptyset. \tag{12.1}$$

In other words, there exists an element $\nu \in \operatorname{supp} \sigma$, *such that* $\mu(\nu)$ *belongs to* $\operatorname{supp} \partial\sigma$.

Proof. We start by proving (1). Set $A := \operatorname{supp} \sigma$, and assume $\sigma \neq 0$, or, equivalently, $A \neq \emptyset$. Let $\mathcal{L}$ be the subcomplex of the simplicial complex $\mathcal{K}$ obtained by taking the union of the closures of all the simplices in A. In other words, the simplicial complex $\mathcal{L}$ consists of all the simplices contained in one of the simplices in A. This a pure simplicial complex of dimension d, and its maximal simplices are indexed by the set A.

By our construction, for each $\beta \in A$, the simplex $\mu(\beta)$ is defined and has dimension $d-1$. Consider a matching λ on $\mathcal{L}$ which matches each d-simplex $\beta \in A$ with the $(d-1)$-simplex $\mu(\beta)$. Since this is a restriction of the acyclic matching μ, the matching λ is itself a well-defined acyclic matching. The unmatched part of $\mathcal{L}$ forms a subcomplex $\widetilde{\mathcal{L}}$, whose dimension is strictly less than d. We now use Theorem 10.9 to conclude that the simplicial complex $\mathcal{L}$ can be collapsed to its subcomplex $\widetilde{\mathcal{L}}$. In particular, $H_d(\mathcal{L}) = 0$, which, since d is the top dimension of $\mathcal{L}$, means that every d-dimensional cycle of $\mathcal{L}$ must in fact be 0. Since $\partial_d \sigma = 0$ in $\mathcal{L}$ as well, we conclude that $\sigma = 0$.

Let us now show (2). As before, set $A := \operatorname{supp} \sigma$, and now let H be the subgraph of $G_d(\mu)$ induced by the vertex set A. This graph is acyclic, since the original graph $G_d(\mu)$ is acyclic. Let v be any of the sources of H, and set $w := \mu(v)$. By construction, $w \in \mu(\operatorname{supp} \sigma)$. On the other hand, there can be no $u \in A \setminus v$, such that $w \in \operatorname{supp} \partial u$, as otherwise the edge (u, v) would belong to H, contradicting the fact that v is a source. However, if there is no such u, then we must have $v \in \operatorname{supp} \partial \sigma$, and we are done. $\square$

Clearly, (2) implies (1) in Proposition 12.2, as $\partial \sigma = 0$ and $\sigma \neq 0$ being valid at the same time would contradict Equation (12.1).

12.2. The closure map φ

Let us now define a function φ which maps each vertex of $G_d(\mu)$ to a set of vertices of $G_d(\mu)$. Algebraically, we will let φ map each d-simplex of $\mathcal{K}$ to a d-chain of $\mathcal{K}$ with coefficients in $\mathbb{Z}_2$.

Definition 12.3. We define φ recursively on the sets V_i of $G_d(\mu)$, in the order of growing indices. As the base, we set $\varphi(\alpha) := \alpha$, for all $\alpha \in V_0$, that is, when α is a sink.

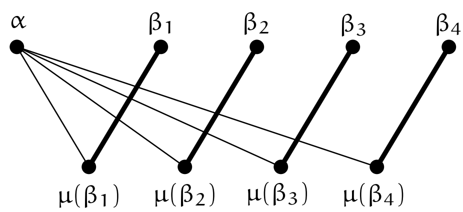

Figure 12.5. The setting defining the function φ in Definition 12.3.

Next, assume that the function φ has already been defined for all vertices in the set $V_0 \cup \cdots \cup V_{i-1}$, for some $1 \leqslant i \leqslant t$, and pick a vertex α in V_i. Let $(\alpha, \beta_1), \dots, (\alpha, \beta_m)$ be the complete list of edges emanating from α. The sets $V_0, \dots, V_t$ were constructed so that all edges point from higher-indexed levels to the lower-indexed ones, so we know that $\beta_1, \dots, \beta_m \in V_0 \cup \cdots \cup V_{i-1}$. We then set

$$\varphi(\alpha) := \alpha + \sum_{i=1}^{m} \varphi(\beta_i), \tag{12.2}$$

which is well-defined by what we just said; see also Figure 12.5.

For the running example from Figures 12.2 to 12.4 we obtain the values

$$\begin{aligned}
\varphi(abe) &= abe,\\
\varphi(ace) &= ace,\\
\varphi(acd) &= acd + ace,\\
\varphi(bce) &= bce + abe,\\
\varphi(bcd) &= bcd + bce + abe,\\
\varphi(abd) &= abd + bcd + bce + abe + acd + ace.
\end{aligned}$$

Let us make a couple of remarks.

Remark 12.4. A criticial simplex σ must be a *source* in graph $G_d(\mu)$. This is because, whenever $G_d(\mu)$ has an edge (α, β), the simplex β must be matched by μ, so it cannot be a critical one.

The converse of Remark 12.4 is false. The graph $G_d(\mu)$ may have sources which are not critical.

Remark 12.5. If σ is not critical, then the support of $\varphi(\sigma)$ does not contain any critical simplices. Otherwise, if σ is critical, the support of $\varphi(\sigma)$ must contain a unique critical simplex, namely, σ itself.

Of course, Remark 12.5 follows at once from Remark 12.4.

Once we have defined φ for the d-simplices, we can take the linear extension and obtain a linear map between the chain groups

$$\varphi : C_d(\mathcal{K}; \mathbb{Z}_2) \to C_d(\mathcal{K}; \mathbb{Z}_2).$$

Next, let us look at an alternative way to compute the function φ.

Definition 12.6. Assume $\mathcal{K}$ is an abstract simplicial complex, and μ is an arbitrary matching on the set of simplices of $\mathcal{K}$. Let α and β be arbitrary simplices of $\mathcal{K}$ such that $\dim \alpha = \dim \beta$. A *reaching path* p from α to β is any sequence

$$\alpha \succ \mu(\beta_1) \prec \beta_1 \succ \mu(\beta_2) \prec \beta_2 \succ \cdots \succ \mu(\beta_m) \prec \beta_m = \beta, \tag{12.3}$$

where $m \geqslant 0$, $\beta_1 \neq \beta_2, \ldots, \beta_{m-1} \neq \beta_m$; see Figure 12.6.

Given a reaching path p from α to β, we set $p_\bullet := \alpha$ and $p^\bullet := \beta$.

Proposition 12.7. *Assume we are given a simplicial complex $\mathcal{K}$ and an acyclic matching μ. Let the function φ be given by Definition* 12.3. *Then the following formula is valid:*

$$\varphi(\alpha) = \sum_{p : p_\bullet = \alpha} p^\bullet, \tag{12.4}$$

where the notation means that the sum is taken over all reaching paths p*, such that* $p_\bullet = \alpha$.

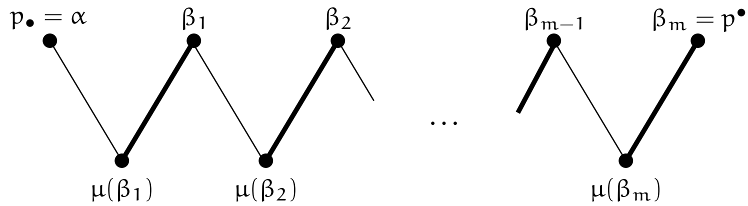

Figure 12.6. An example of a reaching path p.

Proof. Equation (12.4) is obviously valid when α is a sink, i.e., $\alpha \in V_0$, because the only reaching path which starts and terminates in α is the path consisting only of α. We call this path id_α.

We can verify that Equation (12.4) holds for all $\alpha \in V_k$, using induction on k. We already know it is true for $k = 0$, and by the induction hypothesis it is true for V_i, for all $i \leqslant k-1$. By Equation (12.2) we now have

$$\varphi(\alpha) = \alpha + \sum_{i=1}^{m} \varphi(\beta_i) = (\mathrm{id}_\alpha)^\bullet + \sum_{i=1}^{m} \sum_{p:p_\bullet=\beta_i} p^\bullet = \sum_{p:p_\bullet=\alpha} p^\bullet,$$

where the last equality follows from the fact that any non-identity reaching path which starts at the simplex α will first pass through one of the simplices $\beta_1, \dots, \beta_m$. □

Proposition 12.8. *For any* $\sigma \in C_d$ *we have* $\operatorname{supp} \partial(\varphi(\sigma)) \subseteq R_{d-1} \cup M^{\uparrow}_{d-1}$.

Proof. Assume that on the contrary, $\operatorname{supp} \partial(\varphi(\sigma)) \cap M^{\downarrow}_{d-1} \neq \emptyset$, and pick an element β from that intersection. Applying the boundary operator ∂_d to the equation in Proposition 12.7, we have

$$\partial_d \varphi(\sigma) = \sum_p \partial_d p^\bullet, \tag{12.5}$$

where the sum is taken over all reaching paths p such that $p_\bullet = \sigma$. This means that the coefficient with which β appears on the right-hand side of Equation (12.5) is equal to the parity of the number of reaching paths p, such that

(1) $p_\bullet = \sigma$,

(2) the d-simplex $p^\bullet$ contains the $(d-1)$-simplex β.

Since we have assumed that $\beta \in \operatorname{supp} \partial(\varphi(\sigma))$, we must have an odd number of such paths.

Let A denote the set of all such paths. We now describe a complete matching ν on the set A; see Figure 12.7. For each path $p \in A$ we match

$$p \overset{\nu}{\longleftrightarrow} p\,\mathrm{XOR}\{\beta, \mu(\beta)\}.$$

In other words, if $\beta \in p$, then since p is a reaching path, and β is a $(d-1)$-simplex, we must have $\mu(\beta) \in p$. Furthermore, we must then also have $p^\bullet = \mu(\beta)$, so we can match p with the path obtained from p by deleting both β and $\mu(\beta)$.

On the other hand, assume $\beta \notin p$ and set $\gamma := p^\bullet$. We know that γ contains β. Furthermore, $\mu(\beta) \neq \gamma$, as otherwise, we would have β in p. This means that neither β nor $\mu(\beta) \in p$, and adding β and $\mu(\beta)$ to p is again a reaching path originating at σ.

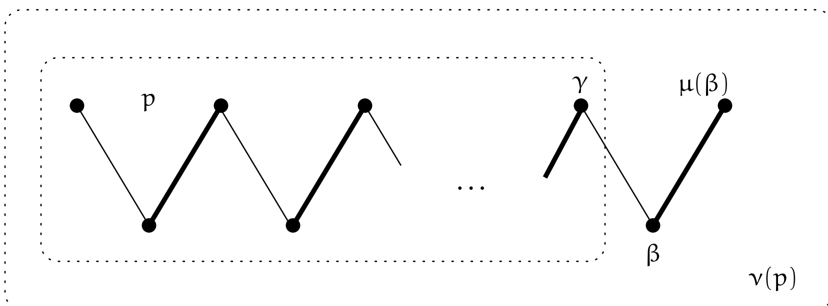

Figure 12.7. Matching ν on the set A.

Thus we have shown that ν is a complete matching on the set A. In particular, the cardinality of the set A is even, contradicting the previous assumption that it was odd. We conclude that $\operatorname{supp} \partial(\varphi(\sigma)) \cap M^{\downarrow}_{d-1}$ is an empty set. $\square$

Also, the converse of Proposition 12.8 is true, in the sense made precise in the next statement.

Proposition 12.9. *Assume* $\operatorname{supp} \sigma \subseteq R_d$ *and* $\operatorname{supp} \alpha \subseteq M^{\uparrow}_d$. *Assume furthermore that* $\operatorname{supp} \partial(\sigma + \alpha) \subseteq R_{d-1} \cup M^{\uparrow}_{d-1}$. *Then* $\varphi(\sigma) = \sigma + \alpha$.

Proof. First, note that by definition of φ, there exists β, such that $\operatorname{supp} \beta \subseteq M^{\uparrow}_d$ and $\varphi(\sigma) = \sigma + \beta$. Let us now consider $\gamma := \varphi(\sigma) + \sigma + \alpha = \beta + \alpha$.[1] Assume $\varphi(\sigma) \neq \sigma + \alpha$, in other words, $\gamma \neq 0$. Obviously, $\operatorname{supp} \gamma \subseteq M^{\uparrow}_d$. On the other hand, by Proposition 12.8 we have $\operatorname{supp} \partial(\varphi(\sigma)) \subseteq R_{d-1} \cup M^{\uparrow}_{d-1}$, and by our assumptions, we have $\operatorname{supp} \partial(\sigma + \alpha) \subseteq R_{d-1} \cup M^{\uparrow}_{d-1}$. It follows that $\operatorname{supp} \partial\gamma \subseteq R_{d-1} \cup M^{\uparrow}_{d-1}$. However, Proposition 12.2(2) tells us that

[1] Remember that the calculations are done mod $\mathbb{Z}_2$.

there must exist $\rho \in \operatorname{supp}\gamma$, such that $\mu(\rho) \in \operatorname{supp}\partial\gamma$. This, of course, is impossible, since $\mu(\rho) \in M^{\downarrow}_{d-1}$, and we have just shown that $\operatorname{supp}\partial\gamma \subseteq R_{d-1} \cup M^{\uparrow}_{d-1}$. We conclude that $\gamma = 0$, and subsequently $\varphi(\sigma) = \sigma + \alpha$. □

Verbally, Propositions 12.8 and 12.9 can be summarized as follows: φ is the closure operator which expands a chain from R with a chain from $M^{\uparrow}$, so that the $M^{\downarrow}$-component of the boundary is annihilated. The operator φ is well-defined since there is a unique way to perform such an extension.

12.3. Homology generators associated to critical cells in an isolated dimension

Assume we are given an abstract simplicial complex $\mathcal{K}$ and an acyclic matching μ. The case we would like to investigate in detail in this subsection is when we only have critical simplices in some chosen dimension d, and no critical simplices in neighboring dimensions. In other words, all the simplices of dimension $d+1$ and $d-1$ are matched by μ. The following statement is the main result of this section.

Theorem 12.10. *Assume, as above, that we are given a simplicial complex $\mathcal{K}$ and an acyclic matching μ, such that, for a certain positive integer d, all the simplices of $\mathcal{K}$ in dimensions $d-1$ and $d+1$ are matched, i.e., $R_{d-1} = R_{d+1} = \emptyset$.*

Assume furthermore that we are given a critical simplex σ of dimension d. Then, with the function φ being defined as above, the value $\varphi(\sigma)$ is a homology class, such that σ is a unique critical d-simplex contained in the support of $\varphi(\sigma)$.

Proof. According to Remark 12.5, we already know that σ is a unique critical d-simplex contained in the support of $\varphi(\sigma)$. Let us now show that $\varphi(\sigma)$ is actually a cycle.

Assume $\partial_d\varphi(\sigma) \neq 0$, and let S denote the support set of $\partial_d\varphi(\sigma) \neq 0$. Set $S_I := S \cap M^{\downarrow}$ and $S_{II} := S \cap M^{\uparrow}$. Since we have assumed that all $(d-1)$-simplices of $\mathcal{K}$ are matched, $\mu(\tau)$ is well-defined for all $\tau \in S$, so the set S is a disjoint union of the sets S_I and S_{II}. This means that we can write

$$\partial_d\varphi(\sigma) = \tau_I + \tau_{II}, \text{ where } \tau_I = \sum_{\tau \in S_I} \tau \text{ and } \tau_{II} = \sum_{\tau \in S_{II}} \tau.$$

The situation is illustrated in Figure 12.8.

All we need to do is to show that $\tau_I = \tau_{II} = 0$. First, $\tau_I = 0$ by Proposition 12.8. Second, the chain τ_{II} must be a cycle. Indeed, since $\tau_I = 0$, we get $\partial_d\varphi(\sigma) = \tau_I + \tau_{II} = \tau_{II}$. Taking boundary once more, we

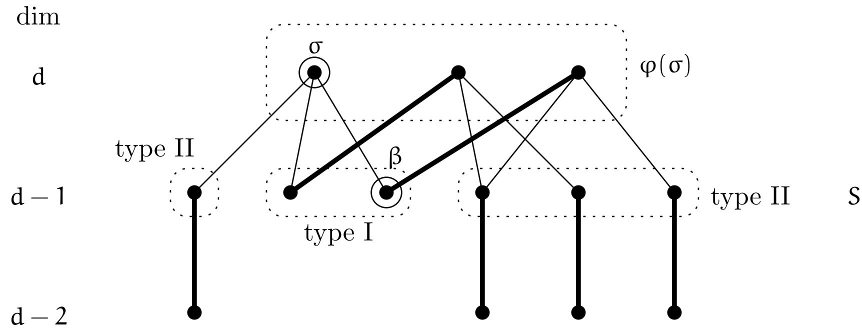

Figure 12.8. Taking the boundary of $\varphi(\alpha)$.

obtain

$$0 = \partial_{d-1}\partial_d\varphi(\sigma) = \partial_{d-1}\tau_{II}.$$

Since $\operatorname{supp}\tau_{II} \subseteq M^{\uparrow}$, we can now use Proposition 12.2(1) to conclude that $\tau_{II} = 0$.

We have shown that $\varphi(\alpha)$ is a cycle. As the last part of the proof, let us see that $\varphi(\alpha)$ represents a non-trivial homology generator. Our main characters are the following disjoint sets:

$$\begin{aligned} M^{\downarrow}_{d+1} &= \{\sigma \in \mathcal{K} \mid \dim\sigma = d+1 \text{ and } \dim\mu(\sigma) = d+2\}, \\ M^{\uparrow}_{d+1} &= \{\sigma \in \mathcal{K} \mid \dim\sigma = d+1 \text{ and } \dim\mu(\sigma) = d\}, \\ M^{\uparrow}_{d+2} &= \{\sigma \in \mathcal{K} \mid \dim\sigma = d+2 \text{ and } \dim\mu(\sigma) = d+1\}, \\ M^{\downarrow}_{d} &= \{\sigma \in \mathcal{K} \mid \dim\sigma = d \text{ and } \dim\mu(\sigma) = d+1\}. \end{aligned}$$

We know that μ matches $M^{\uparrow}_{d+2}$ with $M^{\downarrow}_{d+1}$, it matches $M^{\downarrow}_{d}$ with $M^{\uparrow}_{d+1}$, and the set of $(d+1)$-simplices of $\mathcal{K}$ is a disjoint union of the sets $M^{\downarrow}_{d+1}$ and $M^{\uparrow}_{d+1}$. Let $\mathcal{E}$ denote the subcomplex of the simplicial complex $\mathcal{K}$, defined by

$$\mathcal{E} := \{\sigma \in \mathcal{K} \mid \dim\sigma \leqslant d+1\} \cup M^{\uparrow}_{d+2},$$

and let $\mathcal{F}$ be the subcomplex defined by

$$\mathcal{F} := \mathcal{E} \setminus (M^{\downarrow}_{d} \cup M^{\downarrow}_{d+1} \cup M^{\uparrow}_{d+1} \cup M^{\uparrow}_{d+2}).$$

See Figure 12.9. Since μ is also acyclic on $\mathcal{E}$, Theorem 10.9 implies that there exists a sequence of simplicial collapses from $\mathcal{E}$ to $\mathcal{F}$. In particular, the inclusion map $\iota : \mathcal{F} \hookrightarrow \mathcal{E}$ induces isomorphisms of homology groups.

By our construction, the support of the homology class $\varphi(\sigma)$ is disjoint from the set $M^{\downarrow}_{d}$, hence $\varphi(\sigma)$ is a homology class in the simplicial complex $\mathcal{F}$. Now $\dim\mathcal{F} = d$, so $\varphi(\sigma)$ is in the top dimension. It is also not equal to 0, since it contains σ, so it must be a non-trivial homology class in $\mathcal{F}$. Since the inclusion map $\iota : \mathcal{F} \hookrightarrow \mathcal{E}$ induces isomorphisms of homology groups,

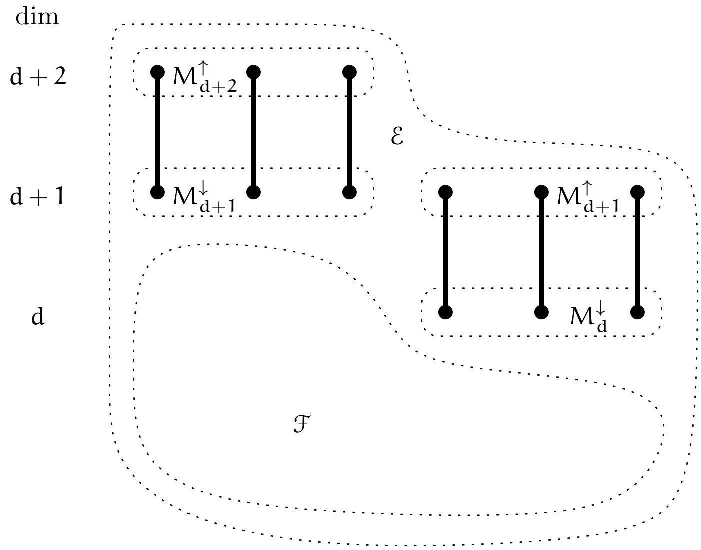

Figure 12.9. The construction of $\mathcal{E}$ and $\mathcal{F}$.

$\varphi(\sigma)$ is also non-trivial in $\mathcal{E}$. It must then also be non-trivial on $\mathcal{K}$, since $\mathcal{E}$ contains the entire $(d+1)$-skeleton of $\mathcal{K}$. Finally, Remark 12.5 tells us that σ is a unique critical d-simplex contained in the support of $\varphi(\sigma)$. □

Remark 12.11. Note that Theorem 12.10 provides a canonical procedure to associate a non-trivial homology generator to an acyclic matching μ and a critical d-simplex, assuming there are no critical simplices in dimensions $d-1$ and $d+1$.

Theorem 12.10 has the following useful corollary.

Corollary 12.12. *Under the same conditions as in Theorem* 12.10*, the homology classes* $[\varphi(\sigma)]$*,* $\sigma \in R_d$ *form a basis for* $H_d(K)$*.*

Proof. We already know that the number of the homology classes $[\varphi(\sigma)]$ is equal to the d-th Betti number, so it is enough to see that they are linearly independent. This can be done by using essentially the same argument as the one in the last part of proof of Theorem 12.10. If a linear combination of $\varphi(\sigma)$'s would represent a trivial homology class, then it would already do so in the subcomplex $\mathcal{F}$. This, however, is impossible, as every non-zero cycle in the top dimension always represents a non-trivial homology class. □

12.4. Sample applications

Let us now see how this knowledge can be applied in practice.

12.4.1. The d-skeleton of an n-simplex. In this example, we return to the situation dealt with in Proposition 11.6. We have seen there that using discrete Morse theory it is easy to show that the d-skeleton of an n-simplex is homotopy equivalent to $\binom{n}{d+1}$ copies of d-dimensional spheres. We can now use Theorem 12.10 to describe explicit homology generators.

Pick a critical d-simplex σ indexed by $0 \leqslant x_0 < \cdots < x_d \leqslant n-1$. The boundary $(d-1)$-simplices of σ are all given by

$$\tau_i := \{x_0, \dots, \widehat{x}_i, \dots, x_d\}, \quad \text{for} \quad 0 \leqslant i \leqslant d.$$

Using the matching μ used in the proof of Proposition 11.6, we obtain

$$\mu(\tau_i) = \{x_0, \dots, \widehat{x}_i, \dots, x_d\} \cup \{n\}.$$

Switching to the graph $G_d(\mu)$, we see that all the vertices $\mu(\tau_0), \dots, \mu(\tau_d)$ are sinks, hence $\varphi(\mu(\tau_i)) = \mu(\tau_i)$, for all $0 \leqslant i \leqslant d$. By Theorem 12.10 we then conclude that the canonical homology class associated to the d-simplex σ under the acyclic matching μ is given by

$$\varphi(\sigma) = \sigma + \mu(\tau_0) + \cdots + \mu(\tau_d).$$

A handy way to write this cycle is to say $\varphi(\sigma) = \partial_{d+1}(\sigma \cup \{n\})$, where the boundary is taken in the full ambient n-simplex.

12.4.2. The complex of disconnected graphs. Consider the simplicial complex Disc_n of disconnected graphs on n vertices, $n \geqslant 3$. In this case each critical $(n-3)$-simplex σ is obtained from a recursive tree T_σ by deleting the edge $(1,2)$. One can see that

$$\varphi(\sigma) = \partial_{n-2} T_\sigma. \tag{12.6}$$

Indeed, for each edge $e \in \sigma$, the graph $\tau_e := \sigma \backslash e$ has 3 connected components, so $\mu(\tau_e)$ is well-defined, and

$$\mu(\tau_e) = (\sigma \setminus e) \cup (1,2).$$

Then, for each $(n-4)$-simplex ν, which is contained in $\mu(\tau_e)$, such that $\nu \neq \tau_e$, we have $(1,2) \in \nu$, hence $\mu(\nu)$ is well-defined, and $\mu(\nu) = \nu \setminus (1,2)$. So, each $\mu(\tau_e)$ is a sink of the directed graph $G_{n-3}(\mu)$, and Equation (12.6) follows.

12.4.3. Order complex of the partition lattice Π_n. Let us now return to considering the order complex of the lattice of all partitions of the set $[n]$, $n \geqslant 3$. Our notation for this simplicial complex was $\widetilde{\Delta}(\Pi_n)$. This complex has dimension $n-3$.

Let μ be the acyclic matching which we have described in Subsection 11.5.2. For this matching, the critical simplices of μ all have dimension $n-3$, so our framework applies and we may attempt to calculate the canonical homology classes associated to the critical simplices.

Recall that each critical simplex is obtained as follows. Choose a permutation π of $[n]$ of the form $1x_2 \dots x_n$. The $(n-3)$-simplex σ_π is a chain of partitions $\alpha_2 < \dots < \alpha_{n-1}$, such that

$$\alpha_i = 1x_2 \dots x_i | x_{i+1} | \dots | x_n.$$

We would like to describe the canonical homology class $\varphi(\sigma_\pi)$. The standard action of the permutation group S_{n-1}, on the set of elements $x_2 \dots x_n$, induces an action on the chains in Π_n. This action is transitive on the set of critical simplices. Hence, if we want to calculate $\varphi(\sigma_\pi)$ in general, it is enough to do that for the ordered tuple $(x_2, \dots, x_n) = (2, \dots, n)$.

So let π be fixed to be the identity permutation. Consider the subposet Q of Π_n consisting of all permutations λ of the form

$$\lambda = 1 \dots i_1 \,|\, i_1 + 1 \dots i_2 \,|\, \dots \,|\, i_t + 1 \dots n,$$

for some $1 \leqslant i_1 < \dots < i_t < n$. It is easy to see that Q is a Boolean algebra on $n-1$ elements. Indeed, each element in Q is simply determined by where the separators $|$ are inserted in the ordered sequence $1, \dots, n$. There are $n-1$ potential positions where the separator can be inserted and any choice of positions is allowed. The order relation is given by inclusion of the sets of positions, and it corresponds precisely to the refinement order in the partition lattice.

Now, the order complex of a Boolean algebra is homeomorphic to a sphere, in fact, it is a barycentric subdivision of the boundary of a simplex. Taking the sum of all the top-dimensional simplices in the order complex of Q yields the fundamental homology class γ_π of the corresponding sphere. It can be shown, see Exercise (5), that γ_π is precisely the canonical homology class associated to the acyclic matching μ and the critical simplex σ_π. When π varies across the set of all $(n-1)!$ permutations of the type above, we get the canonical homology classes associated to all the critical simplices.

Exercises

(1) Use Corollary 12.12 to describe an explicit homology basis for the independence complexes of the graph G, where
  (1) G is a path graph;
  (2) G is an arbitrary tree;
  (3) G is a cycle.

(2) Describe an explicit homology basis for the complex of independent sets of a matroid.

(3) Describe an explicit homology basis for $\mathrm{Flag}_{\mathbf{k}}^{\mathbf{d}}$. The latter complex was defined in Exercise (5) of Chapter 11.

(4) Investigate the computational complexity of finding the homology basis using Corollary 12.12.

(5) Complete the argument in Subsection 12.4.3.