11

Discrete Morse Theory

11.1 Discrete Morse Theory for Posets

When the set of cells of a CW complex is given by means of a combinatorial enumeration, and the cell attachment maps are not too complicated, for instance if the CW complex in question is regular, it is natural to attempt to use the standard notion of *cellular collapse* to simplify the considered topological space, while preserving its homotopy type.

Since the presentation of the cell complex is combinatorial, once this course of action is taken, it becomes imperative to have a language as well as an appropriate combinatorial machinery for dealing with allowed sequences of collapses. Accordingly, we shall first investigate what happens on the purely combinatorial level of posets, before proceeding to drawing topological conclusions and looking at applications.

11.1.1 Acyclic Matchings in Hasse Diagrams of Posets

Recall from Definition 6.13 that for a generalized simplicial complex Δ, *a simplicial collapse* is simply a removal of interiors of two simplices σ and τ such that

- $\dim \sigma = \dim \tau + 1$;
- the only simplex containing σ is σ itself;
- the only simplices containing τ are σ and τ.

Sometimes such a collapse is called an *elementary collapse*. Note that a simplicial collapse is possible if and only if there exists a simplex τ whose link in Δ consists of a single vertex; the simplex σ is then given by the span of τ and v. For a general CW complex one has to take care of some additional technicalities; see Definition 11.12.

In any case, we see that the combinatorial encoding of a set of collapses is best provided by a matching consisting of a collection of pairs of cells (τ, σ)

such that σ contains τ, and $\dim \sigma = \dim \tau + 1$. Clearly, not every matching of this type can be turned into a *sequence of collapses*. For instance, no allowed sequence of collapses for the simplicial complex in Figure 11.1 can be found in the matching depicted on the right of that figure.

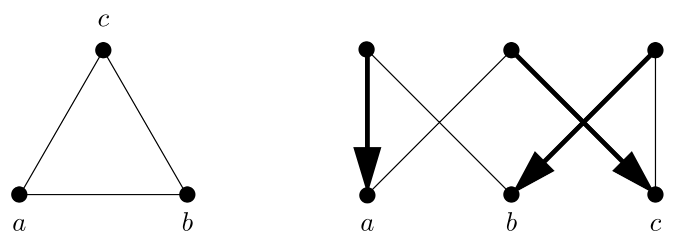

Fig. 11.1. A hollow triangle and a matching in its face poset.

It is easy to see what goes wrong in this example: the prospective collapses are all "hooked up" with each other in a cyclic pattern, which we are unable to break by doing only single collapses. This simple observation leads to the following formalization.

Definition 11.1.
(1) A **partial matching**[1] *in a poset P is a partial matching in the underlying graph of the Hasse diagram of P, i.e., it is a subset $M \subseteq P \times P$ such that*

- *$(a, b) \in M$ implies $b \succ a$;*
- *each $a \in P$ belongs to at most one element in M.*

When $(a, b) \in M$, we write $a = d(b)$ and $b = u(a)$.
(2) A partial matching on P is called **acyclic** *if there does not exist a cycle*

$$b_1 \succ d(b_1) \prec b_2 \succ d(b_2) \prec \cdots \prec b_n \succ d(b_n) \prec b_1, \tag{11.1}$$

with $n \geq 2$, and all $b_i \in P$ being distinct.

A popular way to reformulate condition (2) of Definition 11.1 is the following. Given a poset P, we can orient all edges in the Hasse diagram of P so that they point from the larger element to the smaller one. After that, given a partial matching M, change the orientation of the edges in M to the opposite one. The condition in question now says that the oriented graph obtained in this fashion has no cycles.

We see that Definition 11.1 allows a more general situation than just the collapses that we described above. This makes our situation quite different from the simple homotopy theory considered in Section 6.5. For example, a partial matching consisting of a single pair of simplices $b \succ a$ is always

[1] Also called *discrete vector field*.

acyclic. The reader is invited to intuitively think about such pairs as *internal* collapses. The idea is to remove all the matched elements in some appropriate order, so that the homotopy type of the underlying space is kept intact. We call the unmatched elements, i.e., the elements that will remain, *critical*, and denote the set of critical elements by $C(P, M)$.

The next theorem is the crucial combinatorial fact pertaining to matchings in Hasse diagrams of posets. It characterizes acyclic matchings by means of linear extensions.

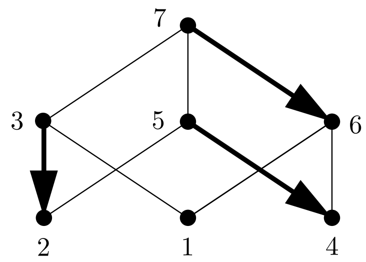

Fig. 11.2. Acyclic matching and the corresponding linear extension.

Theorem 11.2. (Acyclic matchings via linear extensions)
A partial matching on P is acyclic if and only if there exists a linear extension L of P such that the elements a and $u(a)$ follow consecutively in L.

Proof. Assume first that we have a linear extension L satisfying this property, and that at the same time, we have a cycle as in (11.1). Set $a_i = d(b_i)$, for $i = 1, \dots, n$. Then

$$b_{i+1} \succ a_i \Rightarrow a_i <_L b_{i+1} \Rightarrow a_i <_L a_{i+1}$$

(since a_{i+1}, b_{i+1} follow consecutively in L). Thus $a_n >_L a_{n-1} >_L \cdots >_L a_1 >_L a_0 = a_n$, yielding a contradiction.

Assume now that we are given an acyclic matching, and let us define L inductively. Let Q denote the set of elements that are already ordered in L. We start with $Q = \emptyset$. Let W denote the set of minimal elements in $P \setminus Q$. At each step we have one of the following cases.

Case 1. *One of the elements c in W is critical.*

In this case, we simply add c to the order L as the largest element, and proceed with $Q \cup \{c\}$.

Case 2. *All elements in W are matched.*

Consider the subgraph of the underlying graph of the Hasse diagram of $P \setminus Q$ induced by $W \cup u(W)$. Orient its edges as described above, i.e., they should point from the larger element to the smaller one in all cases, except when these two elements are matched, in which case the edge should point from the smaller element to the larger one. Call this oriented graph G.

If there exists an element $a \in W$ such that the only element in $W \cup u(W)$ that is smaller than $u(a)$ is a itself, then we can add elements a and $u(a)$ on top of L and proceed with $Q \cup \{a, u(a)\}$. Otherwise, we see that the outdegree of $u(a)$ in G is positive, for each $a \in W$. On the other hand, the outdegrees in G of all $a \in W$ are equal to 1. Since therefore outdegrees of all vertices in the oriented graph G are positive, we conclude that G must have a cycle, which clearly contradicts the assumption that the considered matching is acyclic.

An example of a linear extension derived from an acyclic matching by this procedure is shown in Figure 11.2. □

11.1.2 Poset Maps with Small Fibers

Next, we would like to characterize acyclic matchings by means of a special class of poset maps.

Definition 11.3. *Given two posets P and Q, a poset map $\varphi : P \to Q$ is said to have* **small fibers** *if for any $q \in Q$, the fiber $\varphi^{-1}(q)$ is either empty or consists of a single element or consists of two comparable elements.*

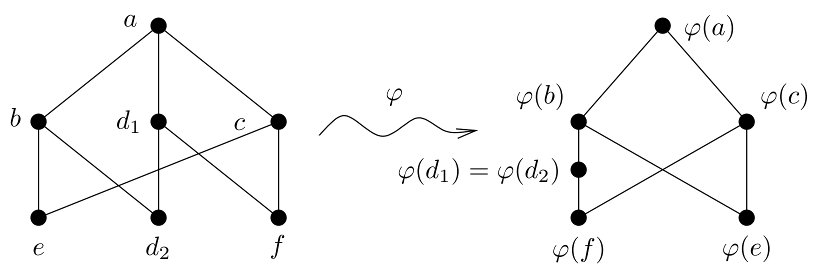

Fig. 11.3. A poset map with small fibers.

An example is shown in Figure 11.3. We remark that since φ is a poset map, if for some $q \in Q$ the fiber $\varphi^{-1}(q)$ consists of two comparable elements, then one of these two elements must actually cover the other one. Therefore, to any given poset map with small fibers $\varphi : P \to Q$ we can associate a partial matching $M(\varphi)$ consisting of all fibers of cardinality 2.

Theorem 11.4. (Acyclic matchings via poset maps with small fibers)
For any poset map with small fibers $\varphi : P \to Q$, the partial matching $M(\varphi)$ is acyclic. Conversely, any acyclic matching on P can be represented as $M(\varphi)$ for some poset map with small fibers φ.

Proof. The fact that $\varphi : P \to Q$ is a poset map implies that the induced matching $M(\varphi)$ is acyclic: for if it were not, there would exist a cycle as in (11.1), and φ would be mapping this cycle to a set of distinct elements $q_1 > q_2 > \cdots > q_t > q_1$ of Q, for some t, yielding a contradiction.

On the other hand, by Theorem 11.2, for any acyclic matching on P there exists a linear extension L of P such that the elements a and $u(a)$ follow consecutively in L. Gluing a with $u(a)$ in this order yields a poset map with small fibers from P to a chain. $\square$

In the proof of Theorem 11.4 we have actually constructed a poset map with small fibers into a chain. These maps are especially important, and we give them a separate name.

Definition 11.5. *A poset map with small fibers $\varphi : P \to Q$ is called a* **collapsing order** *if φ is surjective as a set map, and Q is a chain.*

Given an acyclic matching M, we say that a collapsing order φ is a *collapsing order for M* if it satisfies $M(\varphi) = M$. The etymology of this terminology is fairly clear: the chain Q gives us the order in which it is allowed to perform the prescribed collapses.

11.1.3 Universal Object Associated to an Acyclic Matching

It turns out that for any poset P and any acyclic matching on P, there exists a universal object: a poset whose linear extensions enumerate all allowed collapsing orders.

Definition 11.6. *Let P be a poset, and let M be an acyclic matching on P. We define $U(P, M)$ to be the poset whose set of elements is $M \cup C(P, M)$, and whose partial order is the transitive closure of the elementary relations given by $S_1 \leq_U S_2$, for $S_1, S_2 \in U(P, M)$ if and only if $x \leq y$, for some $x \in S_1$, $y \in S_2$.*

Note that in the formulation of Definition 11.6 we think of elements of M as subsets of P of cardinality 2, while we think of elements of $C(P, M)$ as subsets of P of cardinality 1. One can loosely say that Definition 11.6 states that $U(P, M)$ is obtained from P by gluing each matched pair together to form a single element, with the new partial order induced by the partial order of P in a natural way. See Figure 11.4 for an example.

Of course, the first natural question is whether this new order is actually well-defined. The next proposition answers that question and also explains in what sense $U(P, M)$ is a universal object.

Theorem 11.7. (Universality of $U(P, M)$)
For any poset P and for any acyclic matching M on P, we have:

(1) the partial order on $U(P, M)$ is well-defined;

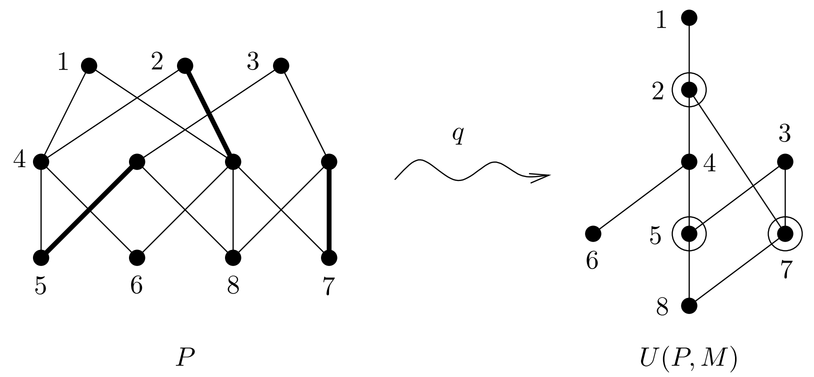

Fig. 11.4. A universal poset associated with an acyclic matching.

(2) the induced quotient map $q : P \to U(P, M)$ *is a poset map with small fibers;*
(3) the linear extensions of $U(P, M)$ *are in 1-to-1 correspondence with collapsing orders for* M*; this correspondence is given by the composition of the quotient map* q *with a linear extension map.*

Proof. To prove (1) we need to check the three axioms of partial orders. The reflexivity is obvious, and the transitivity is automatic, since we have taken the transitive closure. The only property that needs to be proved is antisymmetry. So assume that it does not hold, and take $X, Y \in U(P, M)$ such that $X \leq_U Y$, $Y \leq_U X$, and $X \neq Y$. Choose a sequence

$$X <_U S_1 <_U \cdots <_U S_p <_U Y <_U T_1 <_U \cdots <_U T_q <_U X, \tag{11.2}$$

with the minimal possible p and q. Since p and q are chosen to be minimal, all the sets $S_1, \dots, S_p$ and $T_1, \dots, T_q$ must have cardinality 2.

Let us first deal with the case $p = q = 0$ separately. If $|X| = |Y| = 1$, say $X = \{x\}$, $Y = \{y\}$, then we have $x \leq y$ and $y \leq x$, hence $x = y$, since P itself is a poset. If $|X| = 1$ and $|Y| = 2$, say $X = \{x\}$, $Y = (a, b)$, then $b > x$ and $x > a$, since $x \neq b$, $x \neq a$. This gives $b > x > a$, yielding a contradiction to the assumption that b covers a. By symmetry of (11.2) this argument covers the case $|X| = 2$, $|Y| = 1$ as well, so we can assume that $|X| = |Y| = 2$. In this case $X \leq_U Y \leq_U X$ is a cycle, contradicting the assumption that our matching is acyclic.

From now on, we have $p + q \geq 1$. Assume first that $|X| = |Y| = 1$, say $X = \{x\}$, $Y = \{y\}$. If $p = 0$ and $q = 1$, let $T_1 = (a, b)$, with $b \succ a$. On the one hand, we have $x \leq y$; on the other, $b \geq y$, $x \geq a$. Combining, we get $b \geq y \geq x \geq a$, implying $x = y$, since b covers a. Again by symmetry this takes care of the case $p = 1$ and $q = 0$ as well.

Without loss of generality we may now assume that either $p + q \geq 2$, or $|Y| = 2$ and $p + q \geq 1$. In the first case,

$$S_1 <_U \cdots <_U S_p <_U T_1 <_U \cdots <_U T_q$$

yields a cycle, contradicting the assumption that our matching was acyclic; in the second case such a cycle is given by

$$S_1 <_U \cdots <_U S_p <_U Y <_U T_1 <_U \cdots <_U T_q.$$

Part (2) is straightforward. If $x < y$ in P and $x \in X$, $y \in Y$, for $X, Y \in U(P, M)$, then $X \leq Y$ in $U(P, M)$ by the definition of the partial order on $U(P, M)$, though we may actually get equality. So q is a poset map, and the fibers are small, since we have just proved that $X \leq_U Y$ together with $Y \leq_U X$ implies $X = Y$.

Let us now prove (3). Given a linear extension $l : U(P, M) \to Q$, the composition $l \circ q : P \to Q$ is of course a poset map with small fibers, and it is surjective since both l and q are surjective.

Conversely, assume that $\varphi : P \to Q$ is a collapsing order for M. Since φ is surjective, $\varphi^{-1}(x)$ is nonempty for every $x \in Q$; in fact, we have a bijection between sets $\varphi^{-1}(x)$, for $x \in Q$, and elements of $U(P, M)$. To factor φ through $U(P, M)$, we set $l(q(\varphi^{-1}(x))) := x$, for each $x \in Q$. We have $l \circ q = \varphi$ as set maps. To see that the map l is order-preserving, notice that an elementary relation $S \geq T$, for $S, T \in U(P, M)$, implies that there exist $x \in S$, $y \in T$ such that $x \geq y$, which in turn implies $\varphi(x) \geq \varphi(y)$, since φ is order-preserving, and notice furthermore that all relations $S \geq T$ are just the transitive closures of the elementary ones.

Thus, we get the desired 1-to-1 correspondence between linear extensions of $U(P, M)$, and collapsing orders for M. □

11.1.4 Poset Fibrations and the Patchwork Theorem

Beyond the encoding of all allowed collapsing orders as the set of linear extensions of the universal object $U(P, M)$, viewing the posets with small fibers as the central notion of the combinatorial part of discrete Morse theory is also invaluable for the structural explanation of a standard way to construct acyclic matchings as unions of acyclic matchings on fibers of a poset map.

The following construction generalizes Definition 10.7 of the stack of acyclic categories. Since we will need this only for posets, we satisfy ourselves here with formulating the special case. The generalization to acyclic categories is straightforward.

Definition 11.8. *A* **poset fibration** *is a pair $(B, \mathcal{F})$, where*

- *B is a poset, thought of as the* **base** *of the fibration;*
- *$\mathcal{F} = \{F_x\}_{x \in B}$ is a collection of posets, indexed by the elements of B, thought of as individual* **fibers**.

Associated to such a fibration we have a poset $E(B,\mathcal{F})$, defined as the union $\cup_{x\in B}F_x$, with the order relation given by $\alpha \geq \beta$ if either $\alpha, \beta \in F_x$, and $\alpha \geq \beta$ in F_x, for some $x \in B$, or $\alpha \in F_x$, $\beta \in F_y$, and $x > y$ in B. This is the *total space*.

Furthermore, we have a poset map $p : E(B,\mathcal{F}) \to B$ defined by $p(\alpha) := x$ if $\alpha \in F_x$. In particular, we have $p^{-1}(x) = F_x$, for all $x \in B$. This is the structural *projection map* of the total space to the base space, whose preimages are the fibers.

The notion of poset fibrations satisfies the following universality property.

Theorem 11.9. (Decomposition theorem)
For an arbitrary poset fibration $(B,\mathcal{F})$, where $\mathcal{F} = \{F_x\}_{x\in B}$, and an arbitrary poset P, there is a 1-to-1 correspondence between

- *poset maps $\varphi : P \to E(B,\mathcal{F})$;*
- *pairs $(\psi, \{g_x\}_{x\in B})$, where ψ and each g_x's are poset maps $\psi : P \to B$ and $g_x : \psi^{-1}(x) \to F_x$, for each $x \in B$.*

Under this bijection, the fibers of φ are the same as the fibers of the maps g_x.

Proof. One direction of this bijection is trivial: given a poset map $\varphi : P \to E(B,\mathcal{F})$, we obtain the poset map $\psi : P \to B$ by composing φ with the structural projection map $p : E(B,\mathcal{F}) \to B$, and we obtain the poset maps g_x by taking the appropriate restrictions of the map φ.

In the opposite direction, assume that we have a poset map $\psi : P \to B$ and a collection of poset maps $g_x : \psi^{-1}(x) \to F_x$, for all $x \in B$. Define $\varphi : P \to E(B,\mathcal{F})$ by taking the value of the appropriate fiber map:

$$\psi(\alpha) := g_{\varphi(\alpha)}(\alpha),$$

for all $\alpha \in P$. Let us see that this defines a poset map. For $\alpha > \beta$, $\alpha, \beta \in P$, we have $\varphi(\alpha) \geq \varphi(\beta)$, since φ is a poset map. If $\varphi(\alpha) = \varphi(\beta)$, then $g_{\varphi(\alpha)}(\alpha) \geq g_{\varphi(\beta)}(\beta)$, since $g_{\varphi(\alpha)}$ $(= g_{\varphi(\beta)})$ is a poset map. Otherwise, we have $\varphi(\alpha) > \varphi(\beta)$, and hence $g_{\varphi(\alpha)}(\alpha) > g_{\varphi(\beta)}(\beta)$, by the definition of the partial order on the total space $E(B,\mathcal{F})$. $\square$

The decomposition theorem 11.9, is often used as a rationale to construct an acyclic matching on a poset P in several steps: first map P to some other poset Q, then construct acyclic matchings on the fibers of this map. By the observation above, these acyclic matchings will "patch together" to form an acyclic matching for the whole poset. See Figure 11.5 for an example. For future reference, we summarize this observation in the next theorem.

Theorem 11.10. (Patchwork theorem)
Assume that $\varphi : P \to Q$ is an order-preserving map, and assume that we have acyclic matchings on subposets $\varphi^{-1}(q)$, for all $q \in Q$. Then the union of these matchings is itself an acyclic matching on P.

Proof. The role of the base space here is played by the poset Q, and the fiber maps g_q are given by the acyclic matchings on the subposets $\varphi^{-1}(q)$. The decomposition theorem tells us that there exists a poset map from P to the total space of the corresponding poset fibration, and that the fibers of this map are the same as the fibers of the fiber maps g_q. Since the latter are given by acyclic matchings, we conclude that we have a poset map from P with small fibers that corresponds precisely to the patching of acyclic matchings on the subposets $\varphi^{-1}(q)$, for $q \in Q$. □

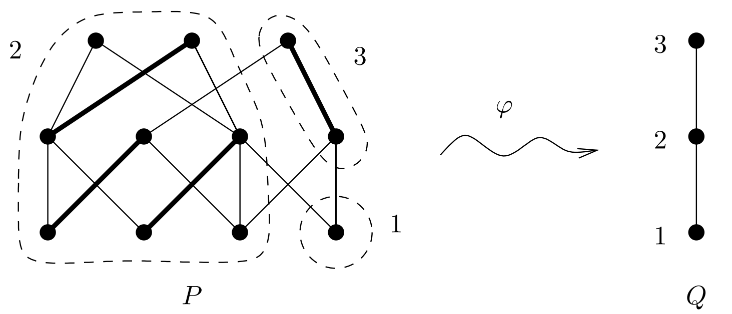

Fig. 11.5. Acyclic matching composed of acyclic matchings on fibers.

We conclude our discussion of poset maps with small fibers by mentioning that this point of view yields a rich class of generalizations. Indeed, any choice of the set of allowed fibers will yield a combinatorial theory that could be interesting to study. One could, for instance, allow any Boolean algebra as a fiber. This would correspond to the theory of all collapses, not just the elementary ones, which we get when considering the small fibers. One can take any other infinite family of posets. One prominent family is that of partition lattices $\{\Pi_n\}_{n=1}^{\infty}$. What happens if we consider all poset maps with partition lattices as fibers?

11.2 Discrete Morse Theory for CW Complexes

11.2.1 Attaching Cells to Homotopy Equivalent Spaces

We shall use the following standard fact of algebraic topology, which we state here with only a sketch of a proof.

Theorem 11.11. *Assume that X_1 and X_1 are two homotopy equivalent topological spaces, and let $h : X_1 \to X_2$ be some homotopy equivalence. Let σ be a cell with attachment maps $f_1 : \partial\sigma \to X_1$ and $f_2 : \partial\sigma \to X_2$ such that $h \circ f_1$ is homotopic to f_2; see Figure 11.6.*

Under these conditions, the space $X_1 \cup_{f_1} \sigma$ *is homotopy equivalent to the space* $X_2 \cup_{f_2} \sigma$.

$$\begin{tikzcd} & \partial\sigma \arrow[dl, "f_1"'] \arrow[dr, "f_2"] & \\ X_1 \arrow[rr, "h"] & & X_2 \end{tikzcd}$$

Fig. 11.6. Attaching a cell to homotopy equivalent spaces.

The homotopy equivalence in Theorem 11.11 can be described by giving an explicit map $f : X_1 \cup_{f_1} \sigma \to X_2 \cup_{f_2} \sigma$. This map is induced by the map h, and by the homotopy $H : \partial\sigma \times I \to X_2$ satisfying $H(\partial\sigma, 0) = f_2$, $H(\partial\sigma, 1) = h \circ f_1$. To describe f, we identify σ with the unit disk D^n, and $\partial\sigma$ with the bounding unit sphere $\mathbb{S}^{n-1}$. Then we set (cf. Corollary 7.12)

$$f(x) := h(x), \qquad \text{for } x \in X_1,$$
$$f(t\mathbf{v}) := \begin{cases} 2t\mathbf{v}, & \text{for } 0 \le t \le 1/2, \quad \mathbf{v} \in \mathbb{S}^{n-1}, \\ H(\mathbf{v}, 2t-1), & \text{for } 1/2 \le t \le 1, \quad \mathbf{v} \in \mathbb{S}^{n-1}. \end{cases}$$

The following two special cases of Theorem 11.11 are often distinguished as being of particular importance:

Case 1. $X_1 = X_2$, *and* $h = \mathrm{id}_{X_1}$.

This is a special case of Proposition 7.11, which is used, for example, in justifying the fact that the homotopy type of a CW complex is uniquely determined even if the cell attachment maps are given only up to homotopy; see Figure 11.7.

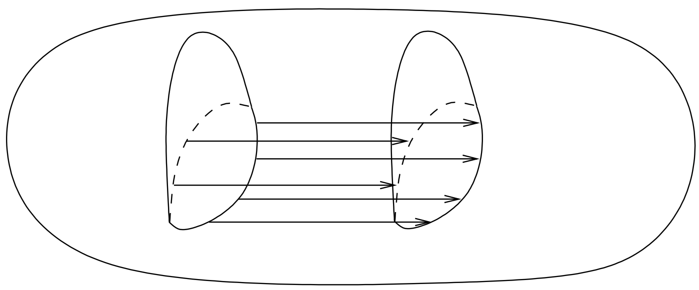

Fig. 11.7. Changing the attachment map by a homotopy.

Case 2. $h \circ f_1 = f_2$.

In fact, if $h \circ f_1 = f_2$, then it is much simpler to describe the homotopy equivalence map $f : X_1 \cup_{f_1} \sigma \to X_2 \cup_{f_2} \sigma$. We may simply set

$$f(x) := \begin{cases} h(x), & \text{for } x \in X_1; \\ x, & \text{for } x \in \operatorname{Int} \sigma. \end{cases} \tag{11.3}$$

11.2.2 The Main Theorem of Discrete Morse Theory for CW Complexes

Intuitively, a cellular collapse is a strong deformation retract that pushes the interior of a maximal cell in, using one of its free boundary cells as the starting point, much like compressing a body made of clay. The cellular collapses can be defined for arbitrary CW complexes.

Definition 11.12. *Let X be a topological space and let Y be a subspace of X. We say that Y is obtained from X by an* **elementary collapse** *if X can be represented as a result of attaching a ball B^n to Y along one of the hemispheres. In other words, if there exists a map $\varphi : B_-^{n-1} \to Y$ such that $X = Y \cup_\varphi B^n$, where B_-^{n-1} denotes one of the closed hemispheres on the boundary of B^n.*

Such a collapse is called **cellular** *if additionally X is a CW complex, and X is a CW complex obtained from Y by attaching two cells: B_+^{n-1} (this is the opposite hemisphere of B_-^{n-1}) and B^n, with φ inducing the necessary attaching maps as above.*

The simplicial collapse defined in Section 6.4 is a special case of Definition 11.12. We are now ready to formulate the central result of this section. For technical convenience, we restrict ourselves to considering cellular collapses in the setting of polyhedral complexes only.

Theorem 11.13.
(Main theorem of discrete Morse theory for CW complexes)
Let Δ be a polyhedral complex, and let M be an acyclic matching on $\mathcal{F}(\Delta)\setminus\{\hat{0}\}$. Let c_i denote the number of critical i-dimensional cells of Δ.

(a) If the critical cells form a subcomplex Δ_c of Δ, then there exists a sequence of cellular collapses leading from Δ to Δ_c.
(b) In general, the space Δ is homotopy equivalent to Δ_c, where Δ_c is a CW complex with c_i cells in dimension i.
(c) There is a natural indexing of cells of Δ_c with the critical cells of Δ such that for any two cells σ and τ of Δ_c satisfying $\dim \sigma = \dim \tau + 1$, the incidence number $[\tau : \sigma]$ is given by

$$[\tau : \sigma] = \sum_c w(c). \tag{11.4}$$

Here the sum is taken over all alternating paths c connecting σ with τ, i.e., over all sequences $c = (\sigma, a_1, u(a_1), \dots, a_t, u(a_t), \tau)$ such that $\sigma \succ a_1$, $u(a_t) \succ \tau$, and $u(a_i) \succ a_{i+1}$, for $i = 1, \dots, a_{t-1}$. For such an alternating path, the quantity $w(c)$ is defined by

$$w(c) := (-1)^t [a_1 : \sigma][\tau : u(a_t)] \prod_{i=1}^{t} [a_i : u(a_i)] \prod_{i=1}^{t-1} [a_{i+1} : u(a_i)], \qquad (11.5)$$

where the incidence numbers in the right-hand side are taken in the complex Δ.

Remark 11.14. The converse of Theorem 11.13(a) is clearly true in the following sense: if Δ_c is a subcomplex of Δ and if there exists a sequence of collapses from Δ to Δ_c, then the matching on the cells of $\Delta \setminus \Delta_c$ induced by this sequence of collapses is acyclic. In particular, a polyhedral complex Δ is collapsible if and only if the poset $\mathcal{F}(\Delta) \setminus \{\hat{0}\}$ allows a complete acyclic matching.

Proof of Theorem 11.13. Take the linear extension L satisfying the conditions of Theorem 11.2.

Proof of (a). Clearly, the linear extension can be chosen so that all the critical cells come first. Hence, if read in decreasing order, L gives a sequence of cellular collapses leading from Δ to Δ_c.

Proof of (b). We perform induction on the cardinality of $\mathcal{F}(\Delta)$. If $|\mathcal{F}(\Delta)| = 1$, the statement is clear. For the induction step, let σ be the last cell in L.

Case 1. *The cell σ is critical.*

Let $\widetilde{\Delta} = \Delta \setminus \operatorname{Int} \sigma$, and let $\varphi : \partial\sigma \to \widetilde{\Delta}$ be the attaching map of σ in Δ. The matching M restricted to $\widetilde{\Delta}$ is again acyclic, and the critical cells are the same, with σ missing. Hence by induction, there exist a CW complex $\widetilde{\Delta}_c$ and a homotopy equivalence $h : \widetilde{\Delta} \to \widetilde{\Delta}_c$.

Consider the composition attaching map $h \circ \varphi : \partial\sigma \to \widetilde{\Delta}_c$; see Figure 11.8. By Theorem 11.11, we conclude that $\widetilde{\Delta} \cup_\varphi \sigma \simeq \widetilde{\Delta}_c \cup_{h\circ\varphi} \sigma$. Note that $\Delta = \widetilde{\Delta} \cup_\varphi \sigma$. The theorem follows by induction if we set $\Delta_c := \widetilde{\Delta}_c \cup_{h\circ\varphi} \sigma$. The new homotopy equivalence map is given by equation (11.3).

$$\begin{tikzcd}
& \partial\sigma \arrow[dl, "\varphi"'] \arrow[dr, "h \circ \varphi"] & \\
\widetilde{\Delta} \arrow[rr, "h"] & & \widetilde{\Delta}_c
\end{tikzcd}$$

Fig. 11.8. Attaching a critical cell.

Case 2. *The cell σ is not critical.*

In this case we must have $(d(\sigma), \sigma) \in M$. Note that $d(\sigma)$ is maximal in $\mathcal{F}(\Delta) \setminus \{\sigma\}$, and let $\widetilde{\Delta} = \Delta \setminus (\operatorname{Int} \sigma \cup \operatorname{Int} d(\sigma))$.

Clearly, removing the pair $(d(\sigma), \sigma)$ is a cellular collapse; in particular, there exists a homotopy equivalence $f : \Delta \to \widetilde{\Delta}$. On the other hand, by the induction assumption, there exist a CW complex $\widetilde{\Delta}_c$ with c_i i-dimensional cells and a homotopy equivalence $\widetilde{f} : \widetilde{\Delta} \to \widetilde{\Delta}_c$. Hence, setting $\Delta_c := \widetilde{\Delta}_c$, we have obtained the desired homotopy equivalence $\widetilde{f} \circ f : \Delta \to \Delta_c$.

Proof of (c). We would like to give an elementary geometric argument. Let σ be some critical cell of Δ of dimension n. Initially, σ was attached along its boundary sphere, but after all the internal collapses, the attachment became more intricate. We would like to envision the attaching map as "a map of the world" drawn on the boundary sphere $\partial\sigma$. When the attaching map changes, we "redraw" this map, usually only locally.

Recall that the collapses can be performed so that the dimension of the collapsing pairs does not increase (the dimension is measured by the one of the two cells that has higher dimension). This means that first all collapses of dimension n can be performed, then all collapses of dimension $n-1$.

We think of this collapsing as a dynamic procedure, and we start by tracing the changes of the attachment map of σ when the collapses of dimension n are executed. Let $(a, u(a))$ be such a collapse. If a was not in the image of the attaching map of σ at this point, then this collapse does not alter the attaching map. If a is in the image of the attaching map of σ at this point, then this collapse alters the attaching map (the map of the world on $\partial\sigma$) as follows: a gets replaced with $\partial u(a) \setminus \operatorname{Int} a$. In a polyhedral complex this says that a gets replaced with the Schlegel diagram of $u(a)$ with respect to a; see Figure 11.9 for an example. This process will continue until all the collapses of dimension n are done.

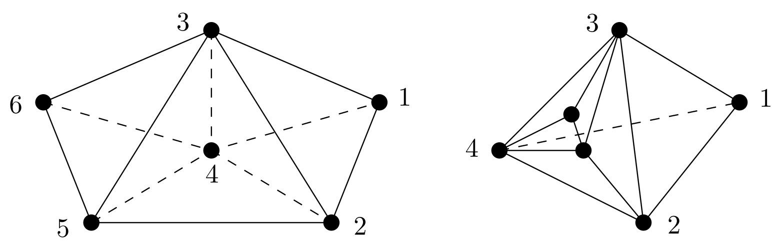

Fig. 11.9. Internal collapses as boundary subdivisions: the result of collapsing $(234, 2345)$ and $(345, 3456)$.

Once it is finished, the only cells of dimension $n-1$ that appear in the image of the attaching map of σ are the critical ones and those that are matched to the cells of dimension $n-2$. The execution of collapses of dimension $n-1$, which follows after that, has a simple effect on the latter ones: they are being internally collapsed, leaving no contribution to the incidence numbers.

This means that the only thing that we need to understand is how often and with which orientations the critical cells of dimension $n-1$ will appear on the boundary sphere $\partial\sigma$. It follows from our iterative procedure above that appearances of a given critical cell τ are in one-to-one correspondence with the alternating paths. Indeed, every replacement of a with $\partial u(a) \setminus \operatorname{Int} a$ corresponds to prolonging the alternating path with the edge up $(a, u(a))$, and then extending it with all possible edges down $(b, u(a))$.

The correctness of (11.4) follows from the following observation, which allows us to trace the evolution of incidence numbers of the cells on $\partial\sigma$. When the cell a with the incidence number ϵ gets replaced by the cells $\partial u(a) \setminus \operatorname{Int} a$, each such cell b gets the incidence number $-\epsilon[a : u(a)][b : u(a)]$. □

It is easy to see that the proof of Theorem 11.13(b) actually works in greater generality: one can take arbitrary CW complexes, at the same time replacing cellular collapses by arbitrary homotopy equivalences. More precisely, we have the following result.

Theorem 11.15. *Let X be a CW complex, and let*

$$F_0(X) \subset F_1(X) \subset \cdots \subset F_t(X) = X$$

be a CW filtration of X such that the subcomplex $F_0(X)$ is just a single vertex, and such that for all $i = 1, \dots, t$, either $F_i(X) \setminus F_{i-1}(X)$ consists of a single cell, or the inclusion map $f_i : F_{i-1}(X) \hookrightarrow F_i(X)$ is a homotopy equivalence.

Then X is homotopy equivalent to a CW complex whose cells are in dimension-preserving bijection with the cells of X, which appear as $F_i(X) \setminus F_{i-1}(X)$.

Proof. If the inclusion map $f_i : F_{i-1}(X) \hookrightarrow F_i(X)$ is a homotopy equivalence, then there exists $g_i : F_i(X) \to F_{i-1}(X)$, which is a homotopy equivalence as well. After this observation, the proof of Theorem 11.13(b) works one-to-one, with critical cells replaced with $F_i(X) \setminus F_{i-1}(X)$ whenever the latter consists of a single cell, and with collapses replaced by such maps g_i. □

11.2.3 Examples

Example 0: Internal collapses on the boundary of a simplex

Our first example is rather simple. Let Δ be the boundary of an n-dimensional simplex. We see that $\mathcal{F}(\Delta) \setminus \{\hat{0}\} = \bar{\mathcal{B}}_{n+1}$. Consider the following matching M on $\bar{\mathcal{B}}_{n+1}$: $(S, S \cup \{1\}) \in M$ for all $S \subset \{2, \dots, n+1\}$. Clearly, this is an acyclic matching, and the only critical simplices are $\sigma = \{1\}$ and $\sigma = \{2, \dots, n+1\}$. Therefore $c_0 = c_{n-1} = 1$, whereas $c_1 = \cdots = c_{n-2} = 0$. It follows that $\Delta \simeq \mathbb{S}^{n-1}$, as is of course expected.

Example 1: The independence complexes of strings and cycles

Our first real application is concerned with the independence complexes of graphs, which were defined in Subsection 9.1.1. Recall that for an arbitrary integer $n \geq 1$, we let L_n denote the graph consisting of n vertices and $n-1$ edges that connect these vertices so as to form a string.

Proposition 11.16. *For any $n \geq 1$, we have*

$$\mathrm{Ind}(L_n) \simeq \begin{cases} \mathbb{S}^{k-1}, & \text{if } n = 3k; \\ \mathrm{pt}, & \text{if } n = 3k+1; \\ \mathbb{S}^{k}, & \text{if } n = 3k+2. \end{cases}$$

Proof. Assume that the vertices of L_n are labeled 1 through n in the same sequence as they occur along the string. Let k denote the maximal integer such that $3k \leq n$. Furthermore, let C be a chain with $k+1$ elements labeled as follows:

$$c_3 > c_6 > \cdots > c_{3k} > c_r.$$

We define a map $\varphi : \mathcal{F}(\mathrm{Ind}\,(L_n)) \to C$ by the following rule. The simplices that contain the vertex labeled 3 get mapped to c_3; the simplices that do not contain the vertex labeled 3, but contain the vertex labeled 6 get mapped to c_6; the simplices that do not contain the vertices labeled 3 and 6, but contain the vertex labeled 9 get mapped to c_9; and so on. Finally, the simplices that contain none of the vertices labeled 3, 6, ..., $3k$ all get mapped to c_r (the index r stands here for *the rest*).

Clearly, the map φ is order-preserving, since if one takes a larger simplex, it will have more vertices, and the only way its image may change is to go up when a new element from the set $\{3, 6, \ldots, 3k\}$ is added and is smaller than the previously smallest one.

Let us now define acyclic matchings on the preimages of C under the map φ. We split our argument into three cases.

Case 1. First we consider the preimage $\varphi^{-1}(c_3)$. For any simplex $\sigma \in \varphi^{-1}(c_3)$ we have $3 \in \sigma$ and hence $2 \notin \sigma$. It follows that pairing $\sigma \leftrightarrow \sigma \oplus \{1\}$ provides a matching that is well-defined and is obviously acyclic.

Case 2. Next, we consider the preimages $\varphi^{-1}(c_6)$ through $\varphi^{-1}(c_{3k})$. Let t be an integer such that $2 \leq t \leq k$. The preimage $\varphi^{-1}(c_{3t})$ consists of all simplices σ such that $3, 6, \ldots, 3t-3 \notin \sigma$, while $3t \in \sigma$. In particular, we have $3t-1 \notin \sigma$. This means that the pairing $\sigma \leftrightarrow \sigma \oplus \{3t-2\}$ provides a well-defined matching, which is acyclic.

Case 3. Finally, we consider the preimage $\varphi^{-1}(c_r)$. We consider three subcases.

If $n = 3k+1$, then this preimage is a face poset with a cone with apex in n; in particular, the pairing $\sigma \leftrightarrow \sigma \oplus \{n\}$ provides an acyclic matching with one critical cell $\sigma = \{n\}$, which has dimension 0. By Theorem 11.10 we can conclude that $\mathrm{Ind}\,(L_{3k+1})$ is collapsible.

If $n = 3k$, we see that $X = \varphi^{-1}(c_r)$ is a face poset of the boundary of a k-dimensional cross-polytope, which is the same as the k-fold join of $\mathbb{S}^0$ with itself. By Theorem 11.10 the matching constructed up to now is acyclic, and it gives us a collapsing sequence leading to X. In particular, this shows that $\mathrm{Ind}\,(L_{3k})$ is homotopy equivalent to $\mathbb{S}^{k-1}$.

If $n = 3k + 2$, we see that $X = \varphi^{-1}(c_r)$ is a face poset of the boundary of a $(k + 1)$-dimensional cross-polytope, since this time around we have the k-fold join of $\mathbb{S}^0$ with itself. The rest is just the same, and we conclude that $\mathrm{Ind}\,(L_{3k+2})$ is homotopy equivalent to $\mathbb{S}^k$. $\square$

Note that the proof of Proposition 11.16 actually yields a stronger statement: instead of contractibility, we actually get a collapsibility, and instead of a mere homotopy equivalence to a sphere of some dimension, we get a sequence of collapses, leading to an explicit sphere, sitting inside $\mathrm{Ind}\,(L_n)$ as a subcomplex.

For an arbitrary integer $n \geq 2$, we let C_n denote the cycle with n vertices. These vertices are labeled $1, \dots, n$, with arithmetic operations on labels performed modulo n. The homotopy type of the independence complexes of cycles allows an easy description as well.

Proposition 11.17. *For any $n \geq 2$, we have*

$$\mathrm{Ind}(C_n) \simeq \begin{cases} \mathbb{S}^{k-1} \vee \mathbb{S}^{k-1}, & \text{if } n = 3k; \\ \mathbb{S}^{k-1}, & \text{if } n = 3k \pm 1. \end{cases}$$

Proof. Let k denote the maximal integer such that $3k \leq n + 1$, and let the chain C be defined in the same way as in Proposition 11.16. Furthermore, let $\varphi : \mathcal{F}(\mathrm{Ind}\,(C_n)) \to C$ be the order-preserving map also described by the same rule as the one in Proposition 11.16. Again, we are looking for acyclic matchings on the preimages.

To start with, the matchings on the preimages $\varphi^{-1}(c_6)$ through $\varphi^{-1}(c_{3k})$ defined identically to Proposition 11.16 are again well-defined and acyclic, without any critical cells. The cases of the remaining two preimages are slightly different.

The preimage $\varphi^{-1}(c_3)$ is the same as the face poset of $\mathrm{Ind}\,(L_{n-3})$ with a minimal element added. Taking the acyclic matching for $\mathrm{Ind}\,(L_{n-3})$ and augmenting it by matching the critical 0-cell with the minimal element yields a new acyclic matching. If $n = 3k+1$, this matching has no critical cells at all. Otherwise, i.e., if $n = 3k$, or $n = 3k - 1$, it has one critical cell in dimension $k - 1$.

Finally, we describe an acyclic matching on the preimage $\varphi^{-1}(c_r)$ by considering three cases.

If $n = 3k - 1$, then we know that $3, 6, \dots, 3k - 3, 3k \notin \sigma$, where we recall that with our conventions $3k = 1$. Therefore, we are dealing with a face poset of a cone with apex in 2, and hence pairing $\sigma \leftrightarrow \sigma \oplus \{2\}$ gives a well-defined acyclic matching with one critical cell $\{2\}$ in dimension 0.

If $n = 3k$, then we again have a face poset of the join of k copies of $\mathbb{S}^0$. Denote the sets of vertices of these k copies of $\mathbb{S}^0$ by $\{x_1, y_1\}, \dots, \{x_k, y_k\}$. Consider the pairing $\sigma \leftrightarrow \sigma \oplus \{x_i\}$, where i is the minimal index such that $y_i \notin \sigma$. This is a well-defined acyclic matching with critical cells $\{x_1\}$ of dimension 0, and $\{y_1, \dots, y_k\}$ of dimension $k-1$.

If $n = 3k+1$, then we have a face poset of $k-1$ copies of $\mathbb{S}^0$ and one copy of $\mathrm{Ind}\,(L_3)$. Denote the sets of vertices of these $k-1$ copies of $\mathbb{S}^0$ by $\{x_1, y_1\}, \dots, \{x_{k-1}, y_{k-1}\}$, and let $\{x_k, y_k, z_k\}$ be the vertices of $\mathrm{Ind}\,(L_3)$, with y_k being the middle vertex. Consider the pairing with the same rule: $\sigma \leftrightarrow \sigma \oplus \{x_i\}$, where i is the minimal index such that $y_i \notin \sigma$. This is a well-defined acyclic matching with critical cells $\{x_1\}$ of dimension 0, and $\{y_1, \dots, y_k\}$ of dimension $k-1$. □

For future reference, let us remark that the proofs of Proposition 11.16 and Proposition 11.17 imply that the inclusion graph homomorphism $i : L_{3k} \hookrightarrow C_{3k+1}$ induces an isomorphism on the homology groups $i_* : H_*(\mathrm{Ind}\,(L_{3k})) \hookrightarrow H_*(\mathrm{Ind}\,(C_{3k+1}))$.

Example 2: The simplicial complex $\Delta(\bar{\Pi}_n)$ is homotopy equivalent to a wedge of $(n-1)!$ spheres of dimension $n-3$

Recall the partition lattice Π_n introduced in Chapter 9.

Theorem 11.18. *For $n \geq 3$, the simplicial complex $\Delta(\bar{\Pi}_n)$ is homotopy equivalent to a wedge of spheres of dimension $n-3$.*

Proof. The statement is obviously true for $n = 3$, so we assume that $n \geq 4$ and proceed by induction. Set $\alpha := (1)(2, 3, \dots, n)$, and let Q to be the interval consisting of all partitions having a singleton block (1) with reversed order, i.e., $Q := [\hat{0}, \alpha]^{\mathrm{op}}_{\Pi_n}$. We define an order-preserving map $\varphi : \mathcal{F}(\Delta(\bar{\Pi}_n)) \to Q$ by the following rule:

a chain c is taken to the minimal element of Q that can be added to c.

One example is shown in Figure 11.10. To analyze this rule, take $c \in \mathcal{F}(\Delta(\bar{\Pi}_n))$, assume $c = (\pi_1 < \pi_2 < \cdots < \pi_t)$, and consider various cases.

Case 1. If $\alpha \geq \pi_t$, then $\varphi(c) = \alpha$.

Case 2. If $\alpha \ngeq \pi_k$ and either $\alpha \geq \pi_{k-1}$ or $k = 1$, then $\varphi(c) = \pi_k \wedge \alpha$.

In words: find the smallest partition π_k in c, where 1 is a part of a nonsingleton block B, and then partition B into (1) and $B \setminus \{1\}$. This also shows that the minimal element in this rule is unique; hence the map φ is well-defined.

To see that φ is order-preserving, it is enough to notice that if the chain is increased, then the minimal possible element of Q, i.e., the maximal possible element of $[\hat{0}, \alpha]_{\Pi_n}$ that can be added to this chain will either remain the same or increase in Q (resp. decrease in $[\hat{0}, \alpha]_{\Pi_n}$).

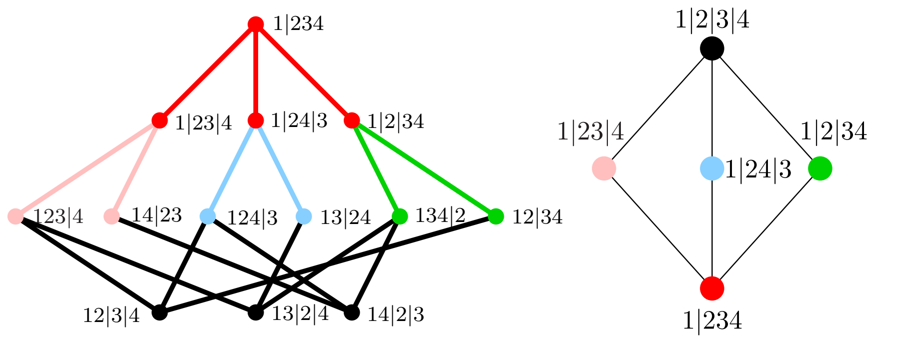

Fig. 11.10. The map φ for $n = 4$.

By Theorem 11.10 it is now sufficient to construct acyclic matchings on the fibers $\varphi^{-1}(x)$. We do this again with case-by-case analysis.

Case 1. Let $S = \varphi^{-1}((1)(2)\dots(n))$. Clearly, the poset S is in fact a disjoint union $S = S_2 \cup \dots \cup S_n$, where S_i is the subposet consisting of all chains containing the element $(1i)(2)\dots(i-1)(i+1)\dots(n)$, for $i = 2, \dots, n$. Each poset S_i is actually a copy of $\mathcal{F}(\Delta(\bar{\Pi}_{n-1})) \cup \{\hat{0}\}$. By induction, there exists an acyclic matching on $\mathcal{F}(\Delta(\bar{\Pi}_{n-1}))$ that has one critical cell in dimension 0 and $(n-2)!$ critical cells in dimension $n-4$.

In the poset $\mathcal{F}(\Delta(\bar{\Pi}_{n-1})) \cup \{\hat{0}\}$ this acyclic matching can be extended to have only the top-dimensional critical element, since the other one is matched with $\hat{0}$. When considered in S_i, these maximal chains consist of $n-2$ elements; hence they correspond to critical simplices of dimension $n-3$ in $\Delta(\bar{\Pi}_n)$.

Case 2. Let $S = \varphi^{-1}(\pi)$, for $\pi \neq (1)(2)\dots(n)$. The matching rule in this case is the following: add π to the chain if it is not there already; otherwise, remove it. Obviously this gives an acyclic matching. The only element that is not matched is the chain consisting only of π: this one would have to be matched with an empty chain, which, by our assumptions, is not there. This corresponds to one critical cell of dimension 0.

Summarizing our considerations, we get $(n-1) \times (n-2)! = (n-1)!$ critical cells of dimension $n-3$ and one critical cell of dimension 0. We may therefore conclude that $\Delta(\bar{\Pi}_n)$ is homotopy equivalent to a wedge of $(n-1)!$ spheres of dimension $n-3$.

The spheres are enumerated by these critical cells of dimension $n-3$. If desired, the recurrence above can be avoided and the chains corresponding to these critical cells can be listed explicitly. These are indexed by permutations of $\{2, \dots, n\}$, where for every such permutation $(i_1, \dots, i_{n-1})$ the corresponding chain c is given by

$$c = (1, i_1)(2)\dots(n) < (1, i_1, i_2)(2)\dots(n) < \dots < (1, i_1, \dots, i_{n-2})(i_{n-1}).$$

The dual cochains of these simplices can also be taken as a basis for the cohomology group $\widetilde{H}^{n-3}(\Delta(\bar{\Pi}_n);\mathbb{Z})$. $\square$

In the next two examples we illustrate how one can check the acyclicity of a partial matching directly, bypassing the patchwork theorem.

Example 3: The generalized simplicial complex $\Delta(\bar{\Pi}_n)/\mathcal{S}_n$ is collapsible

The indexing of the simplices of $\Delta(\bar{\Pi}_n)/\mathcal{S}_n$.
The simplices of $\Delta(\bar{\Pi}_n)$ can be indexed with sequences of the set partitions of $[n]$, where the partitions refine each other. One can think of such a sequence as a forest, where vertices are ordered into levels, each level correspond to a set partition of $[n]$, and each vertex on that level corresponds to a block of this partition.

We can therefore index the simplices of $\Delta(\bar{\Pi}_n)$ with such "leveled forests," where the vertices carry subsets of $[n]$ as labels, and the label of each vertex is equal to the union of the labels of its children; see Figure 11.11 for an example. If one desires, one can also add two artificial levels: one on top, consisting of one vertex having the label $[n]$, and one on the bottom, consisting of n leaves having labels 1 through n; this way we obtain labeled trees. Clearly, the labels of the bottom level determine all other labels.

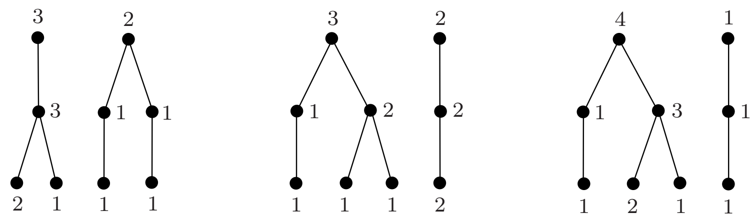

Fig. 11.11. Examples of labeled forests indexing 2-simplices of the generalized simplicial complex $\Delta(\bar{\Pi}_n)/\mathcal{S}_n$.

The symmetric group $\mathcal{S}_n$ acts by permuting the ground set $[n]$. We leave it as an exercise to see that an $\mathcal{S}_n$-orbit of a labeled tree as above consists of all labeled trees with the same cardinalities of the labels on the vertices. As a result of this observation, we can index the simplices of $\Delta(\bar{\Pi}_n)/\mathcal{S}_n$ with the labeled trees as above, with the difference that the labels are positive integers, and labels on each level form a number partition of n (instead of the set partition of $[n]$). For example, the vertices of $\Delta(\bar{\Pi}_n)/\mathcal{S}_n$ are indexed with number partitions of n, edges of $\Delta(\bar{\Pi}_n)/\mathcal{S}_n$ are indexed with the ways two such number partitions can refine each other, and so on.

We also see that, both in the case of $\Delta(\bar{\Pi}_n)$ and in that of $\Delta(\bar{\Pi}_n)/\mathcal{S}_n$, the boundary operator is obtained by deleting entire levels from trees and reconnecting vertices transitively through the deleted level.

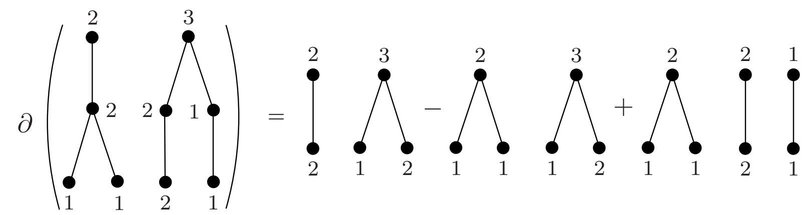

Fig. 11.12. An example of the boundary operator in the generalized simplicial complex $\Delta(\bar{\Pi}_n)/\mathcal{S}_n$.

When λ is a number partition of n, $\lambda = (\lambda_1, \dots, \lambda_t)$, we define $s_2(\lambda) = \sum_{i=1}^{t} \left\lfloor \frac{\lambda_i}{2} \right\rfloor$ and $\mu_2(\lambda) = (2^{s_2(\lambda)}, 1^{n-2\cdot s_2(\lambda)}) = (\underbrace{2,\dots,2}_{s_2(\lambda)}, \underbrace{1,\dots,1}_{n-2\cdot s_2(\lambda)})$. Clearly $\mu_2(\lambda)$ refines λ.

Theorem 11.19. *The generalized simplicial complex $\Delta(\bar{\Pi}_n)/\mathcal{S}_n$ is collapsible, for all $n \geq 3$.*

Proof. Set $X_n := \Delta(\bar{\Pi}_n)/\mathcal{S}_n$. The quotient map $p : \Delta(\bar{\Pi}_n) \to X_n$ is cellular. The vertices of X_n are indexed by the number partitions λ of n, $\lambda \neq (n), (1^n)$. These partitions are partially ordered by refinement. We call the vertices that have the form $(2^a, 1^{n-2a})$, for some a, *special*.

Let us now describe an acyclic matching on $P = \mathcal{F}(X_n)$. Let $F \in P$. If all vertices of F are special, then we put F in the set of critical simplices. If not, define $\lambda(F)$ to be the smallest not special vertex of F. If $\mu_2(\lambda(F))$ is a vertex of F, then we match F with $F \setminus \{\mu_2(\lambda(F))\}$; otherwise, we match F with $F \cup \{\mu_2(\lambda(F))\}$.

To see that this is a valid matching, note that

$$\lambda(F) = \lambda(F \cup \mu_2(\lambda(F))).$$

Next, we show that the obtained matching M is acyclic. Assume that there exists a sequence $\sigma_0, \dots, \sigma_t \in P$ such that all σ_i are different, with the exception $\sigma_0 = \sigma_t$, and such that $u(\sigma_i) \succ \sigma_{i+1}$ for $0 \leq i \leq t-1$. Assume that $u(\sigma_0) = (a_1, \dots, a_\alpha, b_1, \dots, b_\beta)$, where the a_i's are special and b_1 is not. Then $\sigma_0 = (a_1, \dots, a_{\alpha-1}, b_1, \dots, b_\beta)$. Since σ_1 is matched upward, and $u(\sigma_0) \neq u(\sigma_1)$, we have $\sigma_1 = (a_1, \dots, a_\alpha, b_2, \dots, b_\beta)$. We see that the number of special vertices in σ_1 is larger by 1 than in σ_0. Repeating the argument, we see that σ_t has t special vertices more than σ_0; therefore $\sigma_0 \neq \sigma_t$. This leads to the conclusion that M is an acyclic matching.

The critical simplices form a subcomplex of X_n, which we call X_n^C. By Theorem 11.13, there exists a sequence of elementary collapses leading from X_n to X_n^C. Observe that if $A = (a_1, \dots, a_t)$ and $B = (b_1, \dots, b_t)$ are two simplices of $\Delta(\bar{\Pi}_n)$ such that for $i = 1, \dots, t$, both a_i and b_i are of the type $(2^{\alpha_i}, 1^{n-2\alpha_i})$ for some α_i, then there exists $g \in \mathcal{S}_n$ such that $gA = B$, i.e., $p(A) = p(B)$. This implies that X_n^C is a simplex, so we can conclude that X_n is collapsible.

Note that this matching can also be found in a functorial way as follows. Let Q be a chain with $\lfloor n/2 \rfloor$ elements labeled with the numbers $1, \dots, \lfloor n/2 \rfloor$ in reverse order; i.e., 1 labels the maximal element. Define $\varphi : P \to Q$ by mapping the cell $F \in P$ to the maximal number k such that $(2^k, 1^{n-2k})$ is either a vertex of F or can be added to F to form a new cell. This is an order-preserving map, since taking a bigger cell will either keep this number k the same or decrease it.

An acyclic matching on the fiber $\varphi^{-1}(k)$, for $k \in Q$, is simply obtained by adding the vertex $(2^k, 1^{n-2k})$ to cells that do not have it, or removing it from those that do. The only critical cell has dimension 0, and it can be found in the fiber $\varphi^{-1}(k)$, for $k = \lfloor n/2 \rfloor$: it is the vertex $(2^k, 1^{n-2k})$, which cannot be removed, since the empty cell is not being matched. □

Example 4: Bounded sets in a lattice

As we have seen in Chapter 9, there are a number of constructions associating a simplicial complex to a poset (or more generally, to a category); here is yet another one that works for lattices.

Definition 11.20. *Let $\mathcal{L}$ be an arbitrary finite lattice. We define $\mathcal{J}(\mathcal{L})$ to be the simplicial complex whose set of vertices is equal to the set of elements of $\bar{\mathcal{L}}$, and whose simplices are all subsets $S \subseteq \bar{\mathcal{L}}$ that have a nontrivial lower bound, i.e., such that $\bigwedge S \neq \hat{0}$.*

Clearly, the simplicial complex $\mathcal{J}(\mathcal{L})$ contains $\Delta(\bar{\mathcal{L}})$ as a subcomplex. It turns out that much more is true.

Theorem 11.21. *Let $\mathcal{L}$ be an arbitrary finite lattice. Then $\mathcal{J}(\mathcal{L}) \searrow \Delta(\bar{\mathcal{L}})$.*

Proof. As the centerpiece of the argument we define the following partial acyclic matching on $\mathcal{F}(\mathcal{J}(\mathcal{L}))$. Let S be an arbitrary simplex of $\mathcal{J}(\mathcal{L})$. Assume that $\mathcal{F}(\mathcal{J}(\mathcal{L}))[S]$ is not a chain. Set $t := |S|$, and let $\{a_1, a_2, \dots, a_t\}$ be a linear extension of $\mathcal{F}(\mathcal{J}(\mathcal{L}))[S]$, i.e., if $1 \leq i < j \leq t$, then $a_i \not\geq a_j$.

Let $k(S)$ be the maximal index, $1 \leq k(S) \leq t$, such that $a_1 < a_2 < \cdots < a_{k(S)}$, and $a_{k(S)} < a_i$, for all $k(S) + 1 \leq i \leq t$; see Figure 11.13. If S has no minimal element, then we set $k(S) := 0$. Set $a(S) := a_{k(S)+1} \wedge \cdots \wedge a_t$. Since $\mathcal{F}(\mathcal{J}(\mathcal{L}))[S]$ is not a chain, we have $k(S) \leq t - 2$, and hence $a(S)$ is well-defined.

Let Σ be the set of all subsets $S \subseteq \bar{\mathcal{L}}$ such that $\mathcal{F}(\mathcal{J}(\mathcal{L}))[S]$ is not a chain and such that $a(S) \notin S$. For $S \in \Sigma$ define $\mu(S) := S \cup \{a(S)\}$; again see Figure 11.13. Clearly, μ defines a partial matching, and, since for any $S \in \Sigma$ we have $a(\mu(S)) = a(S)$, we see that the set $\mu(\Sigma) \cup \Sigma$ consists of all subsets $S \subseteq \bar{\mathcal{L}}$ such that $\mathcal{F}(\mathcal{J}(\mathcal{L}))[S]$ is not a chain. Consequently, the set of critical elements $\mathcal{C}(\mathcal{F}(\mathcal{J}(\mathcal{L})), \mu)$ consists of all chains $S \in \mathcal{F}(\Delta(\bar{\mathcal{L}}))$.

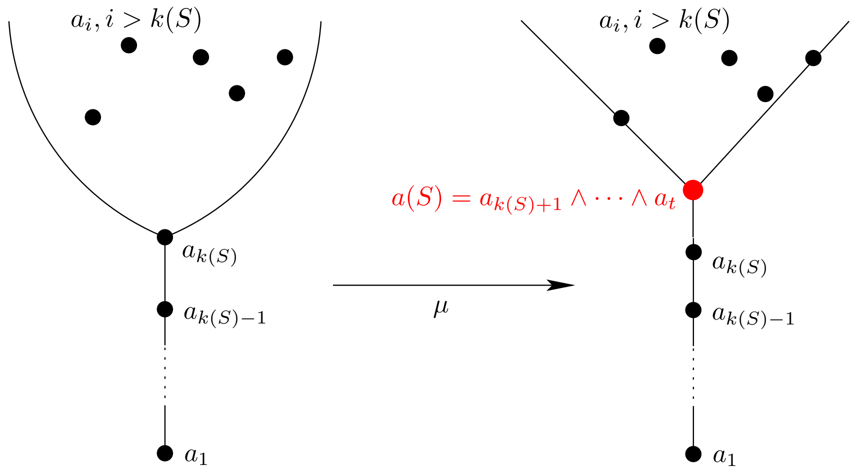

Fig. 11.13. The partial matching μ.

Let us see that the partial matching μ is acyclic. Assume that there exists a sequence $S_1, \dots, S_t \in \Sigma$, where $t \geq 2$, such that $\mu(S_1) \succ S_2$, $\mu(S_2) \succ S_3, \dots,$ $\mu(S_t) \succ S_1$. Let again $\{a_1, a_2, \dots, a_t\}$ be a linear extension of $\mathcal{F}(\mathcal{J}(\mathcal{L}))[S_1]$, as above. By the definition of covering relations, and since $S_2 \neq S_1$, we have $S_2 = \mu(S_1) \setminus \{a_i\}$, for some $1 \leq i \leq t$. If $1 \leq i \leq k(S_1)$, then $a(S_2) = a(S_1)$, which, together with $S_1 = \mu(S_1) \setminus \{a(S_1)\}$, implies $a(S_2) \in S_2$, and hence $S_2 \in \mu(\Sigma)$, giving a contradiction.

Finally, the only option left is that $k(S_1) + 1 \leq i \leq t$, in which case $a(S_2) \geq a(S_1)$, since the join is taken over a set in which each element is larger than $a(S_1)$. If the equality $a(S_2) = a(S_1)$ holds, then $S_2 \in \mu(\Sigma)$, again giving a contradiction. Thus we have shown that a strict inequality must hold: $a(S_2) > a(S_1)$.

Analogously, we can prove that $a(S_{i+1}) > a(S_i)$, for all $1 \leq i \leq t-1$, and that $a(S_1) > a(S_t)$, which, when combined together, yields a contradiction to the assumption that the matching is not acyclic. By Theorem 11.13 we see that the acyclic matching μ provides a sequence of elementary collapses leading from $\mathcal{J}(\mathcal{L})$ to $\Delta(\bar{\mathcal{L}})$. $\square$

More applications will appear in the subsequent sections.

11.3 Algebraic Morse Theory

In this section we give a version of discrete Morse theory that is adapted to the setting of arbitrary free chain complexes.

11.3.1 Acyclic Matchings on Free Chain Complexes and the Morse Complex

Let $\mathcal{R}$ be an arbitrary commutative ring with unit. Recall that a chain complex C_* consisting of $\mathcal{R}$-modules,

$$C_* = \cdots \xrightarrow{\partial_{n+2}} C_{n+1} \xrightarrow{\partial_{n+1}} C_n \xrightarrow{\partial_n} C_{n-1} \xrightarrow{\partial_{n-1}} \ldots,$$

is called *free* if C_n is a finitely generated free $\mathcal{R}$-module for all n. When no confusion can occur, we simply write ∂ instead of ∂_n. We also always require that C_* be bounded on the right.

In order to introduce a combinatorial element into this setting, we need to choose *a basis* (i.e., a set of free generators) Ω_n for each C_n. When this is done, we say that we have chosen a basis $\Omega = \bigcup_n \Omega_n$ for the entire chain complex C_*. We write (C_*, Ω) to denote a chain complex with a basis. A *free chain complex with a basis* is the main object of study of algebraic Morse theory.

Given a free chain complex with a basis (C_*, Ω) and two elements $\alpha \in C_n$ and $b \in \Omega_n$, we denote the coefficient of b in the representation of α as a linear combination of the elements of Ω_n by $\mathbf{k}_\Omega(\alpha, b)$, or, if the basis is clear, simply by $\mathbf{k}(\alpha, b)$. For $x \in C_n$ we write $\dim x = n$. By convention, we set $\mathbf{k}_\Omega(\alpha, b) = 0$ if the dimensions do not match, i.e., if $\dim \alpha \neq \dim b$.

Note that a free chain complex with a basis (C_*, Ω) can be represented as a ranked poset $P(C_*, \Omega)$, with $\mathcal{R}$-weights on the order relations. The elements of rank n correspond to the elements of Ω_n, and the weight of the covering relation $b \succ a$, for $b \in \Omega_n, a \in \Omega_{n-1}$, is simply defined by $w_\Omega(b \succ a) := \mathbf{k}_\Omega(\partial b, a)$. In other words,

$$\partial b = \sum_{b \succ a} w_\Omega(b \succ a) a,$$

for each $b \in \Omega_n$. Again, if the basis is clear, we simply write $w(b \succ a)$.

Definition 11.22. *Let (C_*, Ω) be a free chain complex with a basis. A* **partial matching** $\mathcal{M} \subseteq \Omega \times \Omega$ *on (C_*, Ω) is a partial matching on the covering graph of $P(C_*, \Omega)$ such that if $b \succ a$, and b and a are matched, i.e., if $(a, b) \in M$, then $w(b \succ a)$ is invertible.*

It is important to note that Definition 11.22 is different from the topological one, which was used in Theorem 11.13. In the algebraic setting, the condition that the matched cells form a regular pair (in the CW sense) is replaced by requiring that the covering weights in matched pairs be invertible. However, the notion of *acyclic matching*, which is purely combinatorial, since it is defined on the level of posets, remains the same.

Given such a partial matching $\mathcal{M}$, we denote by $\mathcal{U}_n(\Omega)$ the set of all $b \in \Omega_n$ such that b is matched with some $a \in \Omega_{n-1}$, and analogously, we denote by $\mathcal{D}_n(\Omega)$ the set of all $a \in \Omega_n$ that are matched with some $b \in \Omega_{n+1}$. We let $\mathcal{C}_n(\Omega) := \Omega_n \setminus \{\mathcal{U}_n(\Omega) \cup \mathcal{D}_n(\Omega)\}$ denote the set of critical basis elements of dimension n. Finally, we set $\mathcal{U}(\Omega) := \bigcup_n \mathcal{U}_n(\Omega)$, $\mathcal{D}(\Omega) := \bigcup_n \mathcal{D}_n(\Omega)$, and $\mathcal{C}(\Omega) := \bigcup_n \mathcal{C}_n(\Omega)$.

Given two basis elements $s \in \Omega_n$ and $t \in \Omega_{n-1}$, the *weight* of an alternating path

$$p = (s \succ d(b_1) \prec b_1 \succ d(b_2) \prec b_2 \succ \cdots \succ d(b_n) \prec b_n \succ t), \tag{11.6}$$

where $n \geq 0$ and all $b_i \in \mathcal{U}(\Omega)$ are distinct, is defined to be the quotient

$$w(p) := (-1)^n \frac{w(s \succ d(b_1)) \cdot w(b_1 \succ d(b_2)) \cdots w(b_n \succ t)}{w(b_1 \succ d(b_1)) \cdot w(b_2 \succ d(b_2)) \cdots w(b_n \succ d(b_n))}. \tag{11.7}$$

The reader is invited to compare (11.7) with formula (11.5). Additionally, we shall use the notation $p^\bullet = s$ and $p_\bullet = t$.

Definition 11.23. *Let (C_*, Ω) be a free chain complex with a basis, and let $\mathcal{M}$ be an acyclic matching. The* **Morse complex**

$$\cdots \xrightarrow{\partial^{\mathcal{M}}_{n+2}} C^{\mathcal{M}}_{n+1} \xrightarrow{\partial^{\mathcal{M}}_{n+1}} C^{\mathcal{M}}_{n} \xrightarrow{\partial^{\mathcal{M}}_{n}} C^{\mathcal{M}}_{n-1} \xrightarrow{\partial^{\mathcal{M}}_{n-1}} \cdots,$$

is defined as follows. The $\mathcal{R}$-module $C^{\mathcal{M}}_n$ is freely generated by the elements of $\mathcal{C}_n(\Omega)$. The boundary operator is defined by

$$\partial^{\mathcal{M}}_n(s) = \sum_p w(p) \cdot p_\bullet,$$

for all $s \in \mathcal{C}_n(\Omega)$, where the sum is taken over all alternating paths p satisfying $p^\bullet = s$. Again, if the indexing is clear, we simply write $\partial^{\mathcal{M}}$ instead of $\partial^{\mathcal{M}}_n$.

Given a free chain complex with a basis (C_*, Ω), we can choose a different basis $\widetilde{\Omega}$ by replacing each $a \in \mathcal{D}_n(\Omega)$ by $\tilde{a} = w(u(a) \succ a) \cdot a$, because $w(u(a) \succ a)$ is required to be invertible. Since

$$\mathbf{k}_{\widetilde{\Omega}}(x, \tilde{a}) = \mathbf{k}_\Omega(x, a)/w(u(a) \succ a), \tag{11.8}$$

for any $x \in \Omega_n$, we see that the weights of those alternating paths that do not begin with or end in an element from $\mathcal{D}_n(\Omega)$ remain unaltered, since the quotient $w(x \succ z)/w(y \succ z)$ stays constant as long as $x, y \neq a$. In particular, the Morse complex will not change. On the other hand, by (11.8), $w_{\widetilde{\Omega}}(u(a) \succ$

$a) = 1$, for all $a \in \mathcal{D}(\widetilde{\Omega})$, so the total weight of the alternating path in (11.6) will simply become

$$w_{\widetilde{\Omega}}(p) = (-1)^n w_{\widetilde{\Omega}}(s \succ d(b_1)) \cdot w_{\widetilde{\Omega}}(b_1 \succ d(b_2)) \cdots w_{\widetilde{\Omega}}(b_n \succ t).$$

Because of these observations, we may always replace any given basis of C_* with the basis $\widetilde{\Omega}$ satisfying $w_{\widetilde{\Omega}}(u(a) \succ a) = 1$, for all $a \in \mathcal{D}(\widetilde{\Omega})$.

11.3.2 The Main Theorem of Algebraic Morse Theory

The chain complex $\cdots \longrightarrow 0 \longrightarrow \mathcal{R} \xrightarrow{\text{id}} \mathcal{R} \longrightarrow 0 \longrightarrow \cdots$, where the only nontrivial modules are in the dimensions d and $d-1$, is called an *atom chain complex*, and is denoted by $\texttt{Atom}\,(d)$.

The main theorem of algebraic Morse theory brings to light a certain structure in C_*. Namely, by choosing a different basis, one can represent C_* as a direct sum of two chain complexes, of which one is a direct sum of atom chain complexes, in particular acyclic, and the other one is isomorphic to $C_*^{\mathcal{M}}$. For convenience, the choice of basis can be performed in several steps, one step for each matched pair of the basis elements.

Theorem 11.24.
(Main theorem of discrete Morse theory for free chain complexes)
Assume that we have a free chain complex with a basis (C_*, Ω), and an acyclic matching $\mathcal{M}$. Then C_* decomposes as a direct sum of chain complexes $C_*^{\mathcal{M}} \oplus T_*$, where $T_* \simeq \bigoplus_{(a,b)\in\mathcal{M}} \texttt{Atom}\,(\dim b)$.

It can be advisable to use the example in Subsection 11.3.3 as an illustration for the following proof.

Proof. To start with, let us choose a linear extension L of the partially ordered set $P(C_*, \Omega)$ satisfying the conditions of Theorem 11.2, and let $<_L$ denote the corresponding total order.

Assume first that C_* is bounded; without loss of generality, we can assume that $C_i = 0$ for $i < 0$, and $i > N$. Let $m = |M|$ denote the size of the matching, and let $l = |\Omega| - 2m$ denote the number of critical cells.

We shall now inductively construct a sequence of bases $\Omega^0, \Omega^1, \dots, \Omega^m$ of C_*. More specifically, each basis will be divided into three parts: $\mathcal{C}(\Omega^k) = \{c_1^k, \dots, c_l^k\}$, $\mathcal{D}(\Omega^k) = \{a_1^k, \dots, a_m^k\}$, and $\mathcal{U}(\Omega^k) = \{b_1^k, \dots, b_m^k\}$, such that $a_i^k = d(b_i^k)$, for all $i \in [m]$.

We start with $\Omega^0 = \Omega$ and the initial condition $b_i^0 <_L b_{i+1}^0$, for all $i \in [m-1]$. Since the lower index of $\mathbf{k}_-(-,-)$ and $w_-(- \succ -)$ will be clear from the arguments, we shall omit it to make the formulas more compact.

When constructing the bases, we shall simultaneously prove by induction the following statements:

(i) $C_* = C_*[k] \oplus \mathcal{A}_1^k \oplus \cdots \oplus \mathcal{A}_k^k$, where $C_*[k]$ is the subcomplex of C_* generated by $\Omega^k \setminus \{a_1^k, \dots, a_k^k, b_1^k, \dots, b_k^k\}$, and $\mathcal{A}_i^k$ is isomorphic to $\texttt{Atom}\,(\dim b_i^k)$, for $i \in [k]$;

(ii) for every $x^k \in \{b^k_{k+1}, \dots, b^k_m\} \cup \mathcal{C}(\Omega^k)$, $y^k \in \mathcal{C}(\Omega^k)$, we have $w(x^k \succ y^k) = \sum_p w(p)$, where the sum is restricted to those alternating paths from x^0 to y^0 that use only the pairs (a^0_i, b^0_i), for $i \in [k]$.

Clearly, all of the statements are true for $k = 0$. Assume $k \geq 1$.

Transformation of the basis Ω^{k-1} into the basis Ω^k:
we set

- $a^k_k := \partial b^{k-1}_k$;
- $b^k_k := b^{k-1}_k$;
- $x^k := x^{k-1} - w(x^{k-1} \succ a^{k-1}_k) \cdot b^{k-1}_k$, *for all* $x^{k-1} \in \Omega^{k-1}$, $x \neq a_k, b_k$.

First, we see that Ω^k is a basis. Indeed, assume $b^{k-1}_k \in C_n$. For $i \neq n, n-1$, we have $\Omega^k_i = \Omega^{k-1}_i$; hence by induction, it is a basis. Then Ω^k_{n-1} is obtained from Ω^{k-1}_{n-1} by adding a linear combination of other basis elements to the basis element a^{k-1}_k; hence Ω^k_{n-1} is again a basis. Finally, Ω^k_n is obtained from Ω^{k-1}_n by subtracting multiples of the basis element b^{k-1}_k from the other basis elements; hence it is also a basis.

Next, we investigate how the poset $P(C_*, \Omega^k)$ differs from $P(C_*, \Omega^{k-1})$. If $x \neq b_k$, we have $w(x^k \succ a^k_k) = \mathbf{k}(\partial x^k, a^k_k) = \mathbf{k}(\partial x^k, a^{k-1}_k) = \mathbf{k}(\partial x^{k-1}, a^{k-1}_k) - w(x^{k-1} \succ a^{k-1}_k) \cdot \mathbf{k}(\partial b^{k-1}_k, a^{k-1}_k) = 0$, where the second equality follows from the fact that Ω^k_{n-1} is obtained from Ω^{k-1}_{n-1} by adding a linear combination of other basis elements to the basis element a^{k-1}_k, and the last equality follows from $\mathbf{k}(\partial b^{k-1}_k, a^{k-1}_k) = 1$.

Furthermore, since Ω^k_n is obtained from Ω^{k-1}_n by subtracting multiples of the basis element b^{k-1}_k from the other basis elements, we see that for $x \in \Omega^k_{n+1}$, $y \in \Omega^k_n$, $y \neq b_k$, we have $w(x^k \succ y^k) = w(x^{k-1} \succ y^{k-1})$. Additionally, since the differential of the chain complex squares to 0, we have $0 = \sum_{z^k \in \Omega^k_n} w(x^k \succ z^k) \cdot w(z^k \succ a^k_k) = w(x^k \succ b^k_k) \cdot w(b^k_k \succ a^k_k) = w(x^k \succ b^k_k)$, where the second equality follows from $w(z^k \succ a^k_k) = 0$, for $z \neq b_k$.

We can summarize our findings as follows: all weights in the poset $P(C_*, \Omega^k)$ are the same as in $P(C_*, \Omega^{k-1})$, with the following exceptions:

(1) $w(x^k \succ b^k_k) = 0$, and $w(b^k_k \succ x^k) = 0$, for $x \neq a_k$;

(2) $w(a^k_k \succ x^k) = 0$, and $w(x^k \succ a^k_k) = 0$, for $x \neq b_k$;

(3) $w(x^k \succ y^k) = w(x^{k-1} \succ y^{k-1}) - w(x^{k-1} \succ a^{k-1}_k) \cdot w(b^{k-1}_k \succ y^{k-1})$, for $x \in \Omega^k_n$, $y \in \Omega^k_{n-1}$, $x \neq b_k$, $y \neq a_k$.

In particular, the statement (i) is proved. Furthermore, the following fact can be seen by induction, using (1), (2), and (3):

Fact $(*)$**.** *If $w(x^k \succ y^k) \neq w(x^{k-1} \succ y^{k-1})$, then $b^0_k \geq_L y^0$.*

Indeed, either $y \in \{a_k, b_k\}$ or y is critical or $y = a_{\tilde{k}}$, for $\tilde{k} > k$ such that $w(b^{k-1}_k \succ y^{k-1}) \neq 0$. In the first two cases, $b^0_k \geq_L y^0$ by the construction of L, and the last case is impossible by induction, and again, by the construction of L.

We have $w(b_j^k \succ a_j^k) = w(b_j^{k-1} \succ a_j^{k-1})$, for all j, k. Indeed, this is clear for $j = k$. The case $j < k$ follows by induction, and the case $j > k$ is a consequence of Fact $(*)$.

Next, we see that the partial matching $\mathcal{M}^k := \{(a_i^k, b_i^k) \mid i \in [m]\}$ is acyclic. For $j \leq k$, the poset elements b_j^k, a_j^k are incomparable with the rest; hence they cannot be a part of a cycle. For $i > k$, we have $w(b_j^k \succ a_i^k) = w(b_j^{k-1} \succ a_i^{k-1})$, by Fact $(*)$. Hence by induction, no cycle can be formed by these elements either.

Finally, we trace the boundary operator. Let $x^k \in \{b_{k+1}^k, \dots, b_m^k\} \cup \mathcal{C}(\Omega^k)$, $y^k \in \mathcal{C}(\Omega^k)$. We have $w(x^k \succ y^k) = w(x^{k-1} \succ y^{k-1}) - w(x^{k-1} \succ a_k^{k-1}) w(b_k^{k-1} \succ y^{k-1})$. By induction, the first term counts the contribution of all the alternating paths from x^0 to y^0 that do not use the edges $b_l^0 \succ a_l^0$, for $l \geq k$. The second term contains the additional contribution of the alternating paths from x^0 to y^0 that use the edge $b_k^0 \succ a_k^0$. Observe that if this edge occurs, then by the construction of L, it must be the second edge of the path (counting from x^0), and by Fact $(*)$, we have $w(x^{k-1} \succ a_k^{k-1}) = w(x^0 \succ a_k^0)$. This proves the statement (ii), and therefore concludes the proof of the finite case.

It is now easy to deal with the infinite case, since the basis stabilizes as we proceed through the dimensions, so we may take the union of the stable parts as the new basis for C_*. □

We remark that even if the chain complex C^* is infinite in both directions, one can still define the notions of the acyclic matching and of the Morse complex. Since each particular homology group is determined by a finite excerpt from C^*, we may still conclude that $H_*(C_*) = H_*(C_*^{\mathcal{M}})$.

11.3.3 An Example

Note that the proof of Theorem 11.24 is actually an algorithm. In this subsection we illustrate the workings of this algorithm on a concrete chain complex, namely one associated to some chosen triangulation of the projective plane; see Figure 11.14. For the sake of clarity, we restrict ourselves to $\mathbb{Z}_2$-coefficients. In our figures, a solid line directed from a basis element x down to the basis element y means that ∂x contains y with coefficient 1 when ∂x is decomposed in the *current* basis of the chain group in dimension $\dim y$, that is, the basis consisting of elements that are depicted in the figure in question on the same level as y.

The algorithm starts by picking an extension L satisfying the conditions of Theorem 11.2; this is done in Figure 11.15.

This linear order yields the initial basis Ω^0; see Figure 11.16.

Applying the basis transformation rule to Ω^0 is especially easy, since we are dealing with $\mathbb{Z}_2$-coefficients: all we need to do is to replace all basis elements

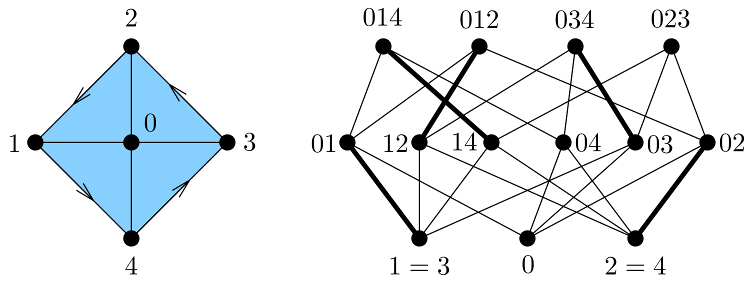

Fig. 11.14. A triangulation of $\mathbb{RP}^2$ and the corresponding face poset.

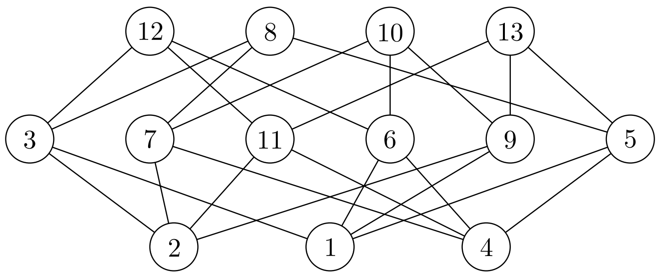

Fig. 11.15. The linear extension L.

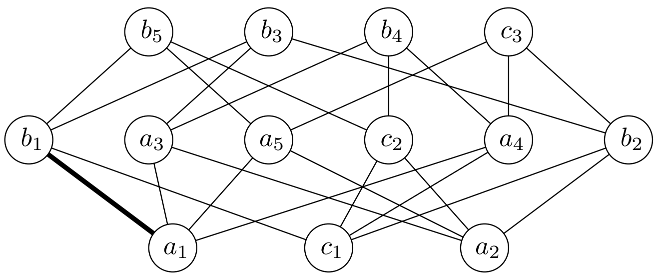

Fig. 11.16. The basis Ω^0.

x containing a_1 in their boundary by $x + b_1$, and recompute the boundaries in the new basis. Again, since we are working over $\mathbb{Z}_2$, the new boundaries of the basis elements in the same dimension as b_1 are simply obtained by taking the symmetric sum (exclusive `or`) of the sets covered by x and by b_1. The analogy with Gaussian elimination is apparent. The resulting basis with corresponding boundaries is shown in Figure 11.17.

We continue in the same way to obtain the bases Ω^2 through Ω^5, as shown in Figures 11.18–11.21. In each figure, the thick line denotes the next collapse. The important thing that one should keep in mind is that at each step, the

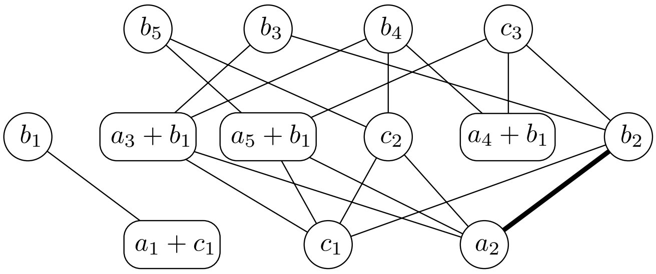

Fig. 11.17. The basis Ω^1.

poset in the figure is given with respect to the new basis, i.e., the boundaries all have to be recalculated accordingly.

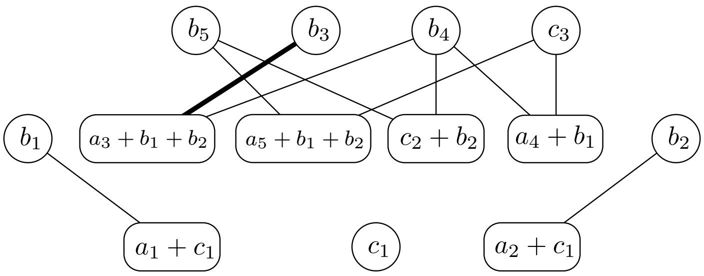

Fig. 11.18. The basis Ω^2.

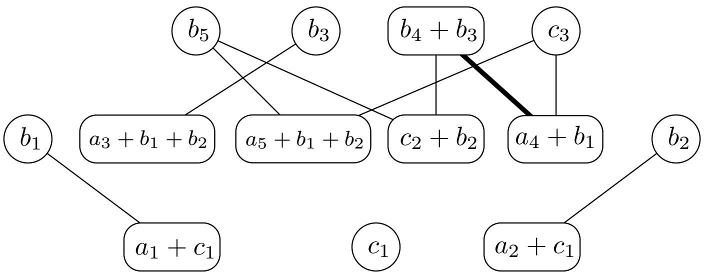

Fig. 11.19. The basis Ω^3.

The final answer is presented in Figure 11.21. It is a good illustration of Theorem 11.24 in this special case. In the new basis, shown in circles and rounded rectangles, our chain complex splits into five atom chain complexes

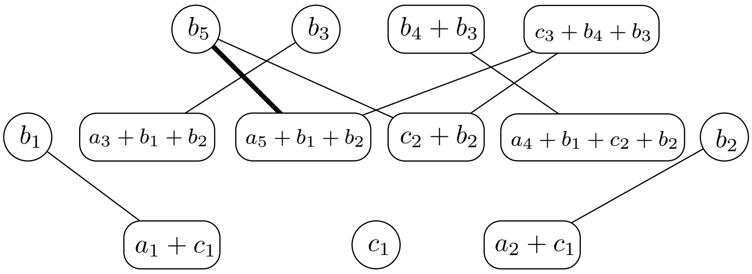

Fig. 11.20. The basis Ω^4.

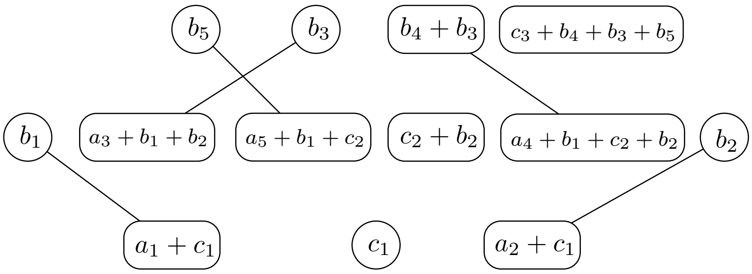

Fig. 11.21. The basis Ω^5.

and the Morse complex. The five atom chain complexes correspond to the initially matched pairs, and have no influence on the homology of the chain complex. The Morse complex is especially simple in this case, since all the differentials $\partial_*^{\mathcal{M}}$ come out to be trivial. This, of course, would have been different if we had worked with integer coefficients. The interested reader is invited to contemplate the latter case.

11.4 Bibliographic Notes

Discrete Morse theory is a tool that was discovered by Forman, whose original article [For98], as well as a later survey [For03], are warmly recommended as excellent sources of background information, as well as some topics that we did not cover in this chapter.

Our main innovation in Section 11.1 is the equivalent reformulation of acyclic matchings in terms of poset maps with small fibers, as well as the introduction of the universal object connected to each acyclic matching. The patchwork theorem 11.10 is a standard tool, used previously by several authors. We think that the terminology of poset fibrations together with the decomposition theorem 11.9 give the patchworking particular clarity.

The material of Sections 11.2.1 and 11.2.2 is quite standard, though our proof of Theorem 11.13(c) is a new ad hoc argument.

We have taken our examples from various sources. Proposition 11.16 and Proposition 11.17 were both originally proved in [Ko99], though in a different way. Theorem 11.18 has been proved in [Bj80], whereas Theorem 11.19 can be found in [Ko00]. Finally, Theorem 11.21 is taken from [Ko06c].

Algebraic Morse theory was discovered independently by several sets of authors. We invite the reader to consult the original sources [JW05, Ko05c, Sk06], providing excellent insight into various aspects of the subject, as well as supplying further applications. Our treatment here follows the algorithmic presentation in [Ko05c].